

Amendments to the Drawings:

Attached hereto are sixty five (65) replacement sheets of drawings, which include Figs. 1 to 69. The attached replacement sheets replace all original drawing sheets. No new matter has been added.

Attachment: sixty five (65) Replacement Sheets

REMARKS

Claims 1 to 7 of the above identified parent case have been canceled, without prejudice. New claims 8 and 15 have been added and are pending herewith for examination in the above referenced continuation application. The amendments made do not add new matter to the application.

In accordance with 37 C.F.R. § 1.121(b)(3), the Substitute Specification (including the Abstract, but without the claims) contains no new matter. The amendments reflected in the Substitute Specification (including Abstract) are to conform the Specification and Abstract to U.S. Patent and Trademark Office rules or to correct informalities. As required by 37 C.F.R. §§ 1.121(b)(3)(iii) and 1.125(b)(2), a Marked Up Version of the Substitute Specification comparing the Specification of record and the Substitute Specification also accompanies this Preliminary Amendment. Approval and entry of the Substitute Specification (including Abstract) is respectfully requested.


The amendments reflected in the Drawings are to correct informalities. Approval and entry of the amendments to the Drawings are respectfully requested.

It is respectfully submitted that the present invention is new, non-obvious, and useful. Prompt consideration and allowance of these claims is therefore respectfully requested.

Respectfully Submitted,

KENYON & KENYON

Dated: 12 June 2006

By 
Michelle M. Carniaux
Reg. No. 36,098

One Broadway
New York, New York 10004
Telephone: (212) 425-7200
Facsimile: (212) 425-5288
CUSTOMER NO. 2664

DATA PROCESSING METHOD AND DEVICE

5 **FIELD OF THE INVENTION**

The present invention relates to methods of operating and optimum use of reconfigurable arrays of data processing elements.

BACKGROUND INFORMATION

- 10 The limitations of conventional processors are becoming more and more evident. The growing importance of stream-based applications makes coarse-grain dynamically reconfigurable architectures an attractive alternative. See, e.g., R. Hartenstein, R. Kress, & H. Reinig, "A new FPGA architecture for word-oriented datapaths," *Proc. FPL'94*, Springer LNCS, September 1994, at 849; E. Waingold et al., "Baring it all to software: Raw
- 15 machines," *IEEE Computer*, September 1997, at 86-93; PACT Corporation, "The XPP Communication System," Technical Report 15 (2000); see generally
- <http://www.pactcorp.com>. They combine the performance of ASICs, which are very risky and expensive (development and mask costs), with the flexibility of traditional processors. See, for example, J. Becker, "Configurable Systems-on-Chip (CSoC)," (Invited Tutorial),
- 20 *Proc. of 9th Proc. of XV Brazilian Symposium on Integrated Circuit, Design (SBCCI 2002)*, (September 2002).

- The datapaths of modern microprocessors reach their limits by using static instruction sets. In spite of the possibilities that exist today in VLSI development, the basic concepts of
- 25 microprocessor architectures are the same as 20 years ago. The main processing unit of modern conventional microprocessors, the datapath, in its actual structure follows the same style guidelines as its predecessors. Although the development of pipelined architectures or superscalar concepts in combination with data and instruction caches increases the
- performance of a modern microprocessor and allows higher frequency rates, the main
- 30 concept of a static datapath remains. Therefore, each operation is a composition of basic instructions that the used processor owns. The benefit of the processor concept lies in the ability of executing strong control dominant application. Data or stream oriented applications are not well suited for this environment. The sequential instruction execution isn't the right target for that kind of application and needs high bandwidth because of permanent

retransmitting of instruction/data from and to memory. This handicap is often eased by use of caches in various stages. A sequential interconnection of filters, which perform data manipulation without writing back the intermediate results would get the right optimisation and reduction of bandwidth. Practically, this kind of chain of filters should be constructed in a logical way and configured during runtime. Existing approaches to extend instruction sets use static modules, not modifiable during runtime.

Customized microprocessors or ASICs are optimized for one special application environment. It is nearly impossible to use the same microprocessor core for another application without loosing the performance gain of this architecture.

A new approach of a flexible and high performance datapath concept is needed, which allows for reconfiguring the functionality and for making this core mainly application independent without losing the performance needed for stream-based applications.

When using such arrays, it is desired to optimise the way a reconfigurable array, it is desirable to optimize the way in which the array is coupled to other units, e.g. to a processor if the array is used as a coprocessor and/or to optimise. It is also desirable to optimize the way in which the array is configured.

The present invention aims at providing improvements over the prior art.

It is to be noted that the disclosure of the present invention does comprise several major parts in its description that all refer to ways of allowing for an optimum use of the array and hence are closely related to each other.

It is also to be noted that the parts do comprise a plurality of figures that the text relates to however without always giving an exact, precise and correct reference. Yet any deviations from correct referencing will be obvious to the average skilled person.

Further, WO 00/49496 discusses a method for execution of a computer program using a processor that includes a configural functional unit capable of executing reconfigurable instructions, which can be redefined at runtime. A problem with conventionable processor architectures exists if a coupling of, for example, sequential processors is needed and/or technologies such as a data-streaming, hyper-threading, multi-threading, multi-tasking,

execution of parts of configurations, etc., are to be a useful way for enhancing performance. Techniques discussed in prior art, such as WO 02/50665 A1, do not allow for a sufficiently efficient way of providing for a data exchange between the ALU of a CPU and the configurable data processing logic cell field, such as an FPGA, DSP, or other such arrangement. In the prior art, the the data exchange is effected via registers. In other words, it is necessary to first write data into a register sequentially, then retrieve them sequentially, and restore them sequentially as well.

SUBSTITUTE SHEET (RULE 26)

Table of Contents

First major part

1	Executive Summary	4
2	Hardware	6
2.1	Code Optimizations	6
2.1.1	Pipelining / Concurrency / Synchronicity	6
2.1.2	Core frequency / Memory Hierarchy	6
2.1.3	Software / Multitasking Operating Systems	7
2.2	Communication Between the RISC Core and the XPP Core	8
2.2.1	Streaming	8
2.2.2	Shared Memory (Main Memory)	8
2.2.3	Shared Memory (IRAM)	8
2.3	State of the XPP Core	9
2.3.1	Limiting Memory Traffic	10
2.4	Context Switches	10
2.4.1	SMT Virtual Processor Switch	11
2.4.2	Interrupt Service Routine	11
2.4.3	Task Switch	11
2.5	A Load Store Architecture	12
2.5.1	A Basic Implementation	14
2.5.2	Implementation Improvements	17
2.5.3	Support for Data Distribution and Data Reorganization	18
2.6	Software 1 Hardware Interface	20

	2.6.1	Explicit Cache.....	20
	2.7	Another Implementation.....	22
	2.7.1	Programming Model Changes.....	24
	2.7.2	Hiding Implementation Details.....	24
5	3	Program Optimizations.....	25
	3.1	Code Analysis.....	25
	3.1.1	Data Flow Analysis.....	25
	3.1.2	Data Dependence Analysis.....	26
10	3.1.3	Interprocedural Alias Analysis.....	28
	3.1.4	Interprocedural Value Range Analysis.....	29
	3.1.5	Alignment Analysis.....	30
	3.2	Code Optimizations.....	31
	3.2.1	General Transformations.....	31
15	3.2.2	Loop Transformations.....	32
	3.2.3	Data Layout Optimizations.....	40
	3.2.4	Example of application of the optimizations.....	41
	4	Compiler Specification for the PACT XPP.....	45
20	4.1	Introduction.....	45
	4.2	Compiler Structure.....	45
	4.2.1	Code Preparation.....	46
	4.2.2	Partitioning.....	46
	4.2.3	RISC Code Generation and Scheduling.....	47
25	4.3	XPP Compiler for Loops.....	47
	4.3.1	Temporal Partitioning.....	48
	4.3.2	Generation of NML Code.....	48
	4.3.3	Mapping Step.....	48
	4.4	XPP Loop Optimizations Driver.....	48
30	4.4.1	Organization of the System.....	48
	4.4.2	Loop Preparation.....	50
	4.4.3	Transformation of the Data Dependence Graph.....	50
	4.4.4	Influence of the Architectural Parameters.....	50
	4.4.5	Optimizations Towards Hardware Improvements.....	56

	4.4.6 Limiting the Execution Time of a Configuration	59
5	Case Studies	60
	5.1 3x3 Edge Detector	60
5	5.1.1 Original Code	60
	5.1.2 Preliminary Transformations	61
	5.1.3 Pre Code Generation Transformations	63
	5.1.4 XPP Code Generation Transformations	64
	5.1.5 Enhancing Parallelism	68
10	5.1.6 Parameterized Function	70
	5.2 FIR Filter	72
	5.2.1 Original Code	72
	5.2.2 First Solution	73
	5.2.3 A More Parallel Solution	77
15	5.2.4 Other Variant	83
	5.3 Matrix Multiplication	84
	5.3.1 Original Code	84
	5.3.2 Preliminary Transformations	84
	5.3.3 XPP Code Generation	90
20	5.4 Viterbi Encoder	93
	5.4.1 Original Code	93
	5.4.2 Interprocedural Optimizations and Scalar Transformations	94
	5.4.3 Initialization	95
	5.4.4 Butterfly Loop	95
25	5.4.5 Re-Normalization	99
	5.4.6 Final Code	103
	5.5 MPEG2 encoder/decoder	107
	5.5.1 Quantization / Inverse Quantization (quant.c)	107
	5.5.2 Inverse Discrete Cosine Transformation (idct.c)	111
30	5.6 Wavelet	123
	5.6.1 Original Code	123
	5.6.2 Optimizing the Whole Loop Nest	124
	5.6.3 Optimizing the Inner Loops	134

Another problem exists if an external access to data is requested in known devices used, inter alia, to implement functions in the configurable data processing logic cell field, DFP, FPGA, etc., that cannot be processed sufficiently on a CPU-integrated ALU. Accordingly, the data processing logic cell field is practically used to allow for user-defined opcodes that can process data more efficiently than is possible on the ALU of the CPU without further support by the data processing logic cell field. In the prior art, the coupling is generally word-based, not block-based. A more efficient data processing, in particular more efficient than possible with a close coupling via registers, is highly desirable.

Another method for the use of logic cell fields that include coarse- and/or fine-granular logic cells and logic cell elements provides for a very loose coupling of such a field to a conventional CPU and/or a CPU-core in embedded systems. In this regard, a conventional sequential program can be executed on the CPU, for example a program written in C, C++, etc., wherein the instantiation or the data stream processing by the fine- and/or coarse-granular data processing logic cell field is effected via that sequential program. However, a problem exists in that for programming said logic cell field, a program not written in C or another sequential high-level language must be provided for the data stream processing. It is desirable to allow for C-programs to run both on a conventional CPU-architecture as well as on the data processing logic cell field operated therewith, in particular, despite the fact that a quasi-sequential program execution should maintain the capability of data-streaming in the data processing logic cell fields, whereas simultaneously the capability exists to operate the CPU in a not too loosely coupled way.

It is already known to provide for sequential data processing within a data processing logic cell field. See, for example, DE 196 51 075, WO 98/26356, DE 196 54 846, WO 98/29952, DE 197 04 728, WO 98/35299, DE 199 26 538, WO 00/77652, and DE 102 12 621. Partial execution is achieved within a single configuration, for example, to reduce the amount of resources needed, to optimize the time of execution, etc. However, this does not lead automatically to allowing a programmer to translate or transfer high-level language code automatically onto a data processing logic cell field as is the case in common. machine models for sequential processes. The compilation, transfer, or translation of a high-level language code onto data processing logic cell fields according to the methods known for models of sequentially executing machines is difficult.

In the prior art, it is further known that configurations that effect different functions on parts of the area respectively can be simultaneously executed on the processing array and that a change of one or some of the configuration(s) without disturbing other configurations is possible at run-time. Methods and hardware-implemented means for the implementation are known to ensure that the execution of partial configurations to be loaded onto the array is possible without deadlock. Reference is made to DE 196 54 593, WO 98/31102, DE 198 07 872, WO 99/44147, DE 199 26538, WO 00/77652, DE 100 28 397, and WO 02/13000. This technology allows in a certain way a certain parallelism and, given certain forms and interrelations of the configurations or partial configurations for a certain way of multitasking/multi-threading, in particular in such a way that the planning, *i.e.*, the scheduling and/or the planning control for time use, can be provided for. Furthermore, from the prior art, time use planning control means and methods are known that, at least under a corresponding interrelation of configurations and/or assignment of configurations to certain tasks and/or threads to configurations and/or sequences of configurations, allow for a multi-tasking and/or multi-threading.

SUMMARY OF THE INVENTION

Embodiments of the present invention may improve upon the prior art with respect to optimization of the way in which a reconfigurable array is coupled to other units and/or the way in which the array is configured.

A way out of limitations of conventional microprocessors may be a dynamic reconfigurable processor datapath extension achieved by integrating traditional static datapaths with the coarse-grain dynamic reconfigurable XPP-architecture (eXtreme Processing Platform).

Embodiments of the present invention introduce a new concept of loosely coupled

5 implementation of the dynamic reconfigurable XPP architecture from PACT Corp. into a static datapath of the SPARC compatible LEON processor. Thus, this approach is different from those where the XPP operates as a completely separate (master) component within one Configurable System-on-Chip (CsoC), together with a processor core, global/local memory topologies, and efficient multi-layer Amba-bus interfaces. See, for example, J. Becker & M.
10 Vorbach, "Architecture, Memory and Interface Technology Integration of an Industrial/Academic Configurable System-on-Chip (CSoC)," *IEEE Computer Society Annual Workshop on VLSI (WVLSI 2003)*, (February 2003). From the programmer's point of view, the extended and adapted datapath may seem like a dynamic configurable instruction set. It can be customized for a specific application and can accelerate the execution enormously.
15 Therefore, the programmer has to create a number of configurations that can be uploaded to the XPP-Array at run time. For example, this configuration can be used like a filter to calculate stream-oriented data. It is also possible to configure more than one function at the same time and use them simultaneously. These embodiments may provide an enormous performance boost and the needed flexibility and power reduction to perform a series of
20 applications very effective.

1——Executive Summary

~~The study is concerned with three objectives:~~

- 1.——~~Proposal of~~Embodiments of the present invention may provide a hardware
25 framework, which enables~~may enable~~ an efficient integration of ~~thea~~ PACT XPP core into a standard RISC processor architecture.
- 2.——~~Proposal of~~Embodiments of the present invention may provide a compiler for ~~thea~~
coupled RISC+XPP hardware. ~~This compiler decides~~ + XPP hardware. The compiler may
30 decide automatically which part of a source code is executed on the RISC processor and which part is executed on the PACT XPP core.
- 3.——~~Presentation of a number of case studies demonstrating which results may be achieved by using the proposed~~In an example embodiment of the present invention, a C Compiler may

be used in cooperation with the proposed hardware framework for the integration of the PACT XPP core and RISC processor.

5 ~~The~~In an example embodiment of the present invention, the proposed hardware framework ~~accelerates~~may accelerate the XPP core in two respects. First, data throughput ~~is~~may be increased by raising the XPP's internal operating frequency into the range of the RISC's frequency. This, however, ~~means that~~may cause the XPP ~~run~~to run into the same pit ~~like~~as all high frequency processors—, *i.e.*, memory accesses may become very slow compared to processor internal computations. ~~This is why the use of a cache is proposed. It eases~~
10 Accordingly, a cache may be provided for use. The cache may ease the memory access problem for a large range of algorithms, which are well suited for an execution on the XPP. The cache, as a second throughput increasing feature ~~requires~~, may require a controller. Hence ~~a~~ A programmable cache controller is ~~introduced, which manages~~may be provided for managing the cache contents and ~~feeds~~feeding the XPP core. It ~~decouples~~may decouple the
15 XPP core ~~computations from the data transfer so that, for instance~~instance, data preload to a specific cache sector ~~takes~~may take place while the XPP is operating on data located in a different cache sector.

20 ~~Another~~A problem ~~emerging~~which may emerge with a coupled RISC+XPP hardware is ~~concerned with~~concerns the RISC's ~~multi-tasking~~multitasking concept. It ~~becomes~~may become necessary to interrupt computations on the XPP in order to perform a task switch. ~~Multitasking is supported by the proposed compiler, as well as by, the proposed hardware.~~
Embodiments of the present invention may provided for hardware and a compiler that supports multitasking. First, each XPP configuration ~~is~~may be considered as an
25 uninterruptible entity. This means that the compiler, which generates the configurations, ~~takes~~may take care that the execution time of any configuration does ~~not~~not exceed a predefined time slice. Second, the cache controller ~~is~~may be concerned with the saving and restoring of the XPP's state after an interrupt. The proposed ~~cache concept~~ minimizesmay minimize the memory traffic for interrupt handling and frequently may even allowallow
30 avoiding memory accesses at all.

FinallyIn an example embodiment of the present invention, the ~~proposed~~ cache concept ismay be based on a simple internal RAM (IRAM) cell structure⁷ allowing for an easy

scalability of the hardware. For instance, extending the XPP cache size, for instance, ~~requires~~may require not much more than the duplication of IRAM cells.

~~The study proposes~~In an embodiment of the present invention, a compiler for a RISC+XPP system. ~~The objective of the compiler is that real- + XPP system may provide for~~
5 compilation for the RISC + XPP system of real world applications, ~~which are~~ written in the C language, ~~can be compiled for a RISC+XPP system~~. The compiler ~~removes~~may remove the necessity of developing NML (Native Mapping Language) code for the XPP by hand. It ~~is~~may be possible, instead, to implement algorithms in the C language or to directly use
10 existing C applications without much adaptation to the XPP system. The ~~proposed~~ compiler ~~includes~~may include the following three major components to perform the compilation process for the XPP:

1. partitioning of the C source code into RISC and XPP parts;₁
2. transformations to optimize the code for the XPP;₁ and
- 15 3. generating of NML code.

~~Finally the~~The generated NML code ~~is~~may be placed and routed for the XPP.

The partitioning component of the compiler ~~decides~~may decide which parts of an application
20 code can be executed on the XPP and which parts are executed on the RISC. Typical candidates for becoming XPP code ~~are~~may be loops with a large number of iterations whose loop bodies are dominated by arithmetic operations. The remaining source code - including the data transfer code - ~~is~~may be compiled for the RISC.

25 The ~~proposed~~ compiler ~~transforms~~may transform the XPP code such that it is optimized for NML code generation. The transformations included in the compiler ~~comprise~~may include a large number of loop transformations as well as general code transformations. Together with data and code analysis the compiler ~~restructures~~may restructure the code so that it ~~fit~~²s fits into the XPP array and so that the final performance ~~exceeds~~may exceed the pure RISC
30 performance. ~~Finally the~~The compiler ~~generates~~may generate NML code from the transformed program. The whole compilation process ~~is~~may be controlled by an optimization driver which selects the optimal order of transformations based on the source code.

~~The~~ Discussed below with respect to embodiments of the present invention are case studies build a major aspect of, the study. ~~The~~ basis of the selection of the examples which is conducted by the guiding principle that each example ~~stands~~ may stand for a set of typical real-world applications. For each example ~~the study demonstrates the work of the proposed compiler. First the code is partitioned.~~ is demonstrated the work of the compiler according to an embodiment of the present invention. For example, first partitioning of the code is discussed. The code transformations, which ~~are~~ may be done by the compiler, are shown and explained. Some examples require minor source code transformations which ~~must~~ may be performed by hand. ~~The study argues that these~~ These transformations ~~are~~ may be either too expensive, or too specific to make sense to be included in the proposed compiler. Dataflow graphs of the transformed codes are constructed for each example, which ~~are~~ may be used by the compiler to generate the NML code. In addition, the XPP resource usages are shown. The case studies demonstrate that a compiler containing the proposed transformations can generate efficient code from: numerical applications for the XPP. This is possible because the compiler ~~relies~~ may rely on the features of the suggested hardware, like the cache controller.

Other embodiments of the present invention pertain to a realization that for data-streaming data-processing, block-based coupling is highly preferable. This is in contrast to a word-based coupling discussed above with respect to the prior art.

Further, embodiments of the present invention provide for the use of time use planning control means, discussed above with respect to their use in the prior art, for configuring and management of configurations for the purpose of scheduling of tasks, threads, and multi- and hyper-threads.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates components of a LEON architecture.

Fig. 2 shows the pipelined datapath structure of the LEON integer unit.

Fig. 3 illustrates components of a typical PAE.

Fig. 4 is a diagram that illustrates an extended datapath according to an example embodiment of the present invention.

5 Fig. 5 illustrates transmission of data to an extended XPP-based datapath by passing the data through an IO-Port, according to an example embodiment of the present invention.

Fig. 6 illustrates an extended LEON instruction pipeline, according to an example embodiment of the present invention.

10 Fig. 7 is a graph that shows that the benefit brought by XPP rises with the number of iDCT blocks computed by it before reconfiguration.

Fig. 8 is a block diagram of an MPEG-4 decoding algorithm, according to an example embodiment of the present invention.

15 Fig. 9 is a block diagram illustrating components of an example embodiment of the present invention, where an XPP core and a RISC core share a memory hierarchy.

20 Fig. 10 shows an IRAM and configuration cache controller data structures and a usage example, according to an example embodiment of the present invention.

Fig. 11 shows an asynchronous pipeline of an XPP, according to an example embodiment of the present invention.

25 Fig. 12 is a diagram that illustrates tasks of an XPP cache controller as states, according to an example embodiment of the present invention.

Fig. 13 shows simultaneous multithreading according to an example embodiment of the present invention.

30 Fig. 14 shows an example of a cache structure according to an example embodiment of the present invention.

Fig. 15 is a control-flow graph of a piece of a program, according to an example embodiment of the present invention.

5 Fig. 16 illustrates a code and diagram of an example of a true dependence with distance 0 on array 'a', according to an example embodiment of the present invention.

Fig. 17 illustrates a code and diagram of an example of an anti-dependence with distance 0 on array 'b', according to an example embodiment of the present invention.

10 Fig. 18 illustrates a code and diagram of an example of an output dependence with distance 0 on array 'a', according to an example embodiment of the present invention.

15 Fig. 19 illustrates a code and diagram of an example of a dependence with direction vector(=,=) between $S1$ and $S2$ and a dependence with direction vector ($=,=,<$) between $S2$ and $S2$, according to an example embodiment of the present invention.

Fig. 20 illustrates a code and diagram of an example of an anti-dependence with distance vector (0,2), according to an example embodiment of the present invention.

20 Fig. 21 is a graph illustrating information of a flow-sensitive alias analysis versus a flow insensitive alias analysis, according to an example embodiment of the present invention.

Fig. 22 is a diagram that illustrates aligned and misaligned memory accesses.

25 Fig. 23 illustrates merging of arrays, according to an example embodiment of the present invention.

Fig. 24 is a flowchart that illustrates a global view of a compiling procedure, according to an example embodiment of the present invention.

30 Fig. 25 is a flowchart that illustrates a detailed architecture and an internal processing of an XPP Compiler.

Fig. 26 is a diagram that illustrates details of XPP loop optimizations, including their organization, according to an example embodiment of the present invention.

5 Fig. 27 is an expression tree of an edge 3x3 inner loop body, according to an example embodiment of the present invention.

Fig. 28 is an expression tree showing the interchanging of operands of commutative add expressions to reduce an overall tree depth, according to an example embodiment of the present invention.

10 Fig. 29 shows a main calculation network of an edge 3x3 configuration, according to an example embodiment of the present invention.

15 Fig. 30 shows a case of synthesized shift registers, according to an example embodiment of the present invention.

Fig. 31 is a data dependency graph relating to a FIR filter, according to an example embodiment of the present invention.

20 Fig. 32 is a dataflow graph that is achieved in an instance where values of x needed for computation of y are kept in registers, according to an example embodiment of the present invention.

25 Fig. 33 is a dataflow graph representing an inner loop with loop unrolling, according to an example embodiment of the present invention.

Fig. 34 is a data dependency graph for matrix multiplication, according to an example embodiment of the present invention.

30 Fig. 35 is a visualization of array access sequences prior to optimization according to an example embodiment of the present invention.

Fig. 36 is a visualization of array access sequences subsequent to optimization according to an example embodiment of the present invention.

Fig. 37 is a dataflow graph of a synthesized configuration and shows matrix multiplication after unroll and jam, according to an example embodiment of the present invention.

- 5 Fig. 38 is a data flow graph corresponding to a butterfly loop, according to an example embodiment of the present invention.

Fig. 39 is a data flow graph showing modifications to code corresponding to the graph of Fig. 38, according to an example embodiment of the present invention.

10

Fig. 40 illustrates a splitting network, according to an example embodiment of the present invention.

- 15 Fig. 41 is a diagram that illustrates how short values are handled, according to an example embodiment of the present invention.

Fig. 42 is a diagram that illustrates how a merge is done, according to an example embodiment of the present invention.

- 20 Fig. 43 illustrates a changing of values of a block row by row before processing of columns.

Fig. 44 illustrates a possible implementation for saturate(val,n) as an NML schematic using two ALUs, according to an example embodiment of the present invention.

- 25 Fig. 45 is a data flow graph for IDCTCOLUMN_CONFIG.

Fig. 46 is a diagram that illustrates use of two counter macros for address generation, according to an example embodiment of the present invention.

- 30 Fig. 47 is a diagram that illustrates an idling of units of a deep pipeline.
Fig. 48 illustrates a loop interchange, according to an example embodiment of the present invention.

Fig. 49 illustrates use of output IRAM of Config A as input IRAM of Config B to bypass a memory interface for bandwidth optimization, according to an example embodiment of the present invention.

- 5 Fig. 50 illustrates block offsets inside tiles generated by a SUIP counter, according to an example embodiment of the present invention.

- Fig. 51 illustrates a difference in efficiency between an instance where there is no data duplication and instance where there is data duplication according to an example embodiment
10 of the present invention.

Fig. 52 illustrates IDCTROW_CONFIG, IDCTCOLUMN_CONFIG, and
REORDER_CONFIG of an example embodiment of the present invention.

- 15 Fig. 53 is a dataflow graph of loop bodies of wavelet after performance of a step of tree balancing, according to an example embodiment of the present invention.

Fig. 54 is a graphical representation of functions for processing data and event packets that can be configured into an RDFP.

20

Figs. 55-69 each illustrates a CDFG according to a respective embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

25

ENVIRONMENT

- Instruction datapaths of modern microprocessors are constrained by certain limitations because they use static instruction sets driven by the traditional von Neumann or Harvard architectural principles. These limitations may be avoided via a dynamic reconfigurable processor datapath extension achieved by integrating traditional static datapaths with the
30 coarse-grain dynamic reconfigurable XPP architecture. Therefore, a loosely asynchronous coupling mechanism of the corresponding instruction datapath or datapath units has been developed and integrated onto a CMOS 0.13 μ m standard cell technology from UMC. In embodiments of the present invention, the SPARC compatible LEON RISC processor may be used, with its static pipelined instruction datapath extended to be configured and personalized

for specific applications. This compiler-compatible instruction set extension allows various and efficient uses, e.g., in streaming application domains like MPEG-4, digital filters, mobile communication modulation, etc. Discussed below is a coupling technique by flexible dual-clock FIFO interfaces that allows asynchronous concurrency of the additionally configured compound instructions, which are integrated into the programming and compilation environment of the LEON processor, and that allows adaption of the frequency of the configured XPP datapath, dependent on actual performance requirements, e.g., for avoiding unneeded cycles and reducing power consumption.

The coupling technique of embodiments of the present invention discussed below combines the flexibility of a general purpose microprocessor with the performance and power consumption of coarse-grain reconfigurable datapath structures, nearly comparable to ASIC performance. Two programming and computing paradigms (control-driven von Neumann and transport-triggered XPP) are unified within one hybrid architecture with the option of two clock domains. The ability to reconfigure the transport-triggered XPP makes the system independent from standards or specific applications. This concept creates potential to develop multi-standard communication devices like software radios by using one extended processor architecture with adapted programming and compilation tools. Thus, new standards can be easily implemented through software updates. The system is scalable during design time through the scalable array-structure of the used XPP extension. This extends the range of suitable applications from products with less multimedia functions to complex high performance systems.

LEON RISC Microprocessor

Embodiments of the present invention may be implemented using a 32-bit SPARC V8 compatible LEON microprocessor. See SPARC International Inc., *The SPARC Architecture Manual, Version 8*, at <http://www.sparc.com>; Jiri Gaisler, *The LEON Processor User's Manual*, at <http://www.gaisler.com>. This microprocessor is a synthesisable, freely available VHDL model which has a load/store architecture and has a five stages pipeline implementation with separated instruction and data caches.

Fig. 1 illustrates components of a LEON architecture. The LEON may be provided with a full implementation of an AMBA 2.0 AHB and APB on-chip bus (1000, 1002), a hardware multiplier and divider, programmable 8/16/32-bit memory controller 1005 for external

PROM, static RAM and SDRAM, and several on-chip peripherals such as timers 1010, UARTs 1012, an interrupt controller 1014, and a 16-bit I/O port 1016. A simple power down mode may be implemented as well.

5 LEON is developed by the European Space Agency (ESA) for future space missions. The performance of LEON is close to an ARM19 series but does not have a memory management unit (MMU) implementation, which limits the use to single memory space applications. Fig. 2 shows the pipelined datapath structure of the LEON integer unit.

10 eXtreme Processing Platform - XPP

Embodiments of the present invention may be implemented using the XPP architecture. Regarding the XPP architecture, see <http://www.pactcorp.com>; "The XPP Communication System," *supra*; and V. Baumgarte et al., "A Self-Reconfigurable Data Processing Architecture," *The 1st Intl. Conference of Engineering of Reconfigurable Systems and*

15 *Algorithms (ERSA '01)*, Las Vegas, NV (June 2001). The XPP architecture is based on a hierarchical array of coarse-grain, adaptive computing elements called *Processing Array Elements (PAEs)* and a *packet-oriented communication network*. The strength of the XPP technology originates from the combination of array processing with unique, powerful run-time reconfiguration mechanisms. Since configuration control is distributed over a

20 *Configuration Manager (CM)* embedded in the array, PAEs can be configured rapidly in parallel while neighboring PAEs are processing data. Entire applications can be configured and run independently on different parts of the array. Reconfiguration may be triggered externally or even by special event signals originating within the array, enabling self-reconfiguring designs. By utilizing protocols implemented in hardware, data and event

25 packets may be used to process, generate, decompose and merge streams of data.

The XPP has some similarities with other coarse-grain reconfigurable architectures like the KressArray (see R. Hartenstein et al., *supra*) or Raw Machines (see E. Waingold et al., *supra*), which are specifically designed for stream-based applications. XPP's main

30 distinguishing features are its automatic packet-handling mechanisms and its sophisticated hierarchical configuration protocols for runtime and self reconfiguration.

Array Structure

A CM may include a state machine and internal RAM for configuration caching. The PAE itself (see top right-hand side of Fig. 3) may include a configuration bus which connects the CM with PAEs and other configurable objects. Horizontal busses may carry data and events. They can be segmented by configurable switch-objects, and can be connected to PAEs and special I/O objects at the periphery of the device.

A PAE is a collection of PAE objects. Fig. 3 illustrates components of a typical PAE, which may include a BREG object (back registers) 1100 and an FREG object (forward registers) 1102, which are used for vertical routing, as well as an ALU object 1104 which performs the actual computations. The ALU 1104 may perform common fixed-point arithmetical and logical operations as well as several special three input opcodes, such as multiply-add, sort, and counters. Events generated by ALU objects depend on ALU results or exceptions, very similar to the state flags of a conventional microprocessor. A counter, e.g., generates a special event only after it has terminated. How these events are used is discussed below. Another PAE object implemented in the XPP is a memory object which can be used in FIFO mode or as RAM for lookup tables, intermediate results, etc. However, any PAE object functionality can be included in the XPP architecture.

Packet Handling and Synchronization

PAE objects, as defined above, may communicate via a packet-oriented network. Two types of packets may be sent through the array: data packets and event packets. Data packets have a uniform bit width specific to the device type. In normal operation mode, PAE objects are self-synchronizing. An operation is performed as soon as all necessary data input packets are available. The results are forwarded as soon as they are available, provided the previous results have been used. Thus, it is possible to map a signal-flow graph directly to ALU objects. Event packets are one bit wide. They transmit state information which controls ALU execution and packet generation.

Configuration

Every PAE stores locally its current configuration state, i.e., if it is part of a configuration or not (states “configured” or “free”). Once a PAE is configured, it changes its state to “configured.” This prevents the CM from reconfiguring a PAE which is still used by another application. The CM caches the configuration data in its internal RAM until the required PAEs become available.

While loading a configuration, all PAEs start to compute their part of the application as soon as they are in state “configured.” Partially configured applications are able to process data without loss of packets. This concurrency of configuration and computation hides configuration latency.

XPP Application Mapping

The NML language, a PACT proprietary structural language with reconfiguration primitives, was developed by PACT to map applications to the XPP array. It gives the programmer direct access to all hardware features.

In NML, configurations consist of modules which are specified as in a structural hardware description language, similar to, for example, structural VHDL. PAE objects are explicitly allocated, optionally placed, and their connections specified. Hierarchical modules allow component reuse, especially for repetitive layouts. Additionally, NML includes statements to support configuration handling. A complete NML application program may include one or more modules, a sequence of initially configured modules, differential changes, and statements which map event signals to configuration and prefetch requests. Thus, configuration handling is an explicit part of the application program.

XPP-based architectures and development tools, such as the PACT XPP Development Suite (XDS) are discussed in detail at <http://www.pactcorp.com>.

LEON Instruction Datapath Extension

LEON and XPP should be able to communicate with each other in a simple and high performance manner. While the XPP is a dataflow orientated device, the LEON is a general purpose processor, suitable for handling control flow. See, for example, *The SPARC Architecture Manual, supra*; Jiri Gaisler, *supra*. Therefore, LEON may be used for system control. To do this, the XPP is integrated into the datapath of the LEON integer unit, which is able to control the XPP. Fig. 4 is a diagram that illustrates this extended datapath.

Due to unpredictable operation time of the XPP algorithm, integration of XPP into LEON data-path is done in a loosely-coupled way. Thus, the XPP array can operate independently of the LEON, which is able to control and reconfigure the XPP during runtime. Since the

configuration of XPP is handled by LEON, the CM 1106 of the XPP is unnecessary and can be left out of the XPP array. The configuration codes are stored in the LEON RAM. LEON transfers the needed configuration from its system RAM into the XPP and creates the needed algorithm on the array.

- 5 To enable a maximum of independence of XPP from LEON, all ports of the XPP - input ports as well as output ports - are buffered using dual clock FIFOs. Dual-clocked FIFOs are implemented into the IO-Ports between LEON and XPP. To transmit data to the extended XPP-based datapath, the data are passed through an IO-Port as shown in Fig. 5. In addition to the FIFO, the IO-Ports include logic to generate handshake signals and an interrupt request
- 10 signal. The IO-Port for receiving data from XPP is similar to Fig. 5 except with a reversed direction of the data signals. This enables XPP to perform completely independently of LEON as long as there are input data available in the input port FIFOs and free space for result data in the output port FIFOs. There are a number of additional features implemented in the LEON pipeline to control the data transfer between LEON and XPP.

15

When LEON tries to write to an IO-Port containing a full FIFO or read from an IO-Port containing an empty FIFO, a trap is generated. This trap can be handled through a trap handler. A further mechanism, e.g., pipeline-holding, may be implemented to allow LEON to hold the pipeline and wait for free FIFO space during XPP write access or wait for a valid

20 FIFO value during XPP read access. When using pipeline-holding, the software developer has to avoid reading from an IO-Port with an empty FIFO while the XPP, or the XPP input IO-Ports, includes no data to produce output. In this case a deadlock will occur requiring a reset of the complete system.

- 25 XPP can generate interrupts for the LEON when trying to read a value from an empty FIFO port or to write a value to a full FIFO port. The occurrence of interrupts indicates that the XPP array cannot process the next step because it has either no input values or it cannot output the result value. The interrupts generated by the XPP are maskable.

- 30 The interface provides information about the FIFOs. LEON can read the number of valid values that are in the FIFO.

Fig. 6 illustrates an extended LEON instruction pipeline. The interface, shown in Fig. 6, to the XPP appears to the LEON as a set of special registers. These XPP registers can be divided into a communication register category and a status register category.

- 5 For data exchange, the XPP communication registers are used. Since XPP provides three different types of communication ports, there are also three types of communication registers. each type is split into an input part and an output part.

Communication Registers

- 10 The data for the process are accessed through XPP data registers. The number of data input and data output ports, as well as the data bit-width depends on the implemented XPP array.

XPP can generate and consume events. Events are one bit signals. The number of input events and output events also depends on the implemented XPP array.

- 15 Configuration of the XPP is done through the XPP configuration register. LEON reads the required configuration value from a file stored in its system RAM and writes it to the XPP configuration register.

20 Status Registers

- There are a number of XPP status registers implemented to control the behavior and get status information of the interface. Switching between the usage of trap handling and pipeline holding can be done in the hold register. An XPP clock register includes a clock frequency ratio between LEON and XPP. By writing to this register, LEON software can set
- 25 the XPP clock relative to the LEON clock. This allows adaptation of the XPP clock frequency to the required XPP performance and consequently allows for influencing the power consumption of the system. Writing zero to the XPP clock register turns off the XPP. There is also a status register for every FIFO including the number of valid values actually available in the FIFO.

- 30 This status registers provide a high degree of flexibility in communication between LEON and XPP and enables different communication modes.

Modes

If there is only one application running on the system at a particular time, software may be developed in pipeline-hold mode. In this instance, LEON initiates data read or write from or to XPP. If there is no value to read or no value to write, LEON pipeline will be stopped until read or write is possible. This can be used to reduce power consumption of the LEON part.

5

In interrupt mode, XPP can influence the LEON program flow. Thus, the IO-Ports generate an interrupt depending on the actual number of values available in the FIFO. The communication between LEON and XPP is via interrupt service routines.

10 Polling mode is a way to access the XPP. Initiated by a timer event, LEON reads XPP ports including data and writes to XPP ports including free FIFO space. Between these phases, LEON can compute other calculations.

It is possible to switch between these strategies anytime within one application.

15

A conventional XPP includes a configuration manager to handle configuration and reconfiguration of the array. However, in combination with the LEON, the configuration manager is dispensable because the configuration as well as any reconfiguration is controlled by the LEON through the XPP configuration register. All XPP configurations used for an application are stored in the LEON's system RAM.

20

Tool and Compiler Integration

To make the new XPP registers accessible through software, the LEON's SPARC 8 instruction set (see *The SPARC Architecture Manual, supra*) is extended by a new subset of instructions. These instructions are based on the SPARC instruction format, but do not conform to the SPARC V8 standard. Corresponding to the SPARC conventions of a load/store architecture, the instruction subset can be divided into two categories. Load/store instructions can exchange data between the LEON memory and the XPP communication registers. The number of cycles per instruction is similar to the standard load/store instructions of the LEON. Read/write instructions are used for communications between LEON registers. Since the LEON register-set is extended by the XPP registers, the read/write instructions are also extended to access XPP registers. Status registers can only be accessed with read/write instructions. Execution of arithmetic instructions on XPP registers are not

25

30

possible. Values have to be written to standard LEON registers before they can be targets of arithmetic operations.

The complete system can still execute any SPARC V8 compatible code. Doing this, the XPP is completely unused.

The LEON is provided with the LECCS cross compiler system (see *LEON/ERC32 Cross Compilation System (LECCS)* at http://www.gaisler.com/cms4_5_3/index.php?option=com_content&task=view&id=62&Itemid=149) under the terms of LGPL. This system includes modified-versions of the binutils 2.11 and gcc 2.95.2. To make the new instruction subset available to software developers, the assembler of the binutils has been extended by a number of instructions according to the implemented instruction subset. The new instructions have the same mnemonic as the regular SPARC V8 load, store, read, and write instructions. Only the new XPP registers have to be used as a source or target operand. Since the modifications of LECCS are straightforward extensions, the cross compiler system is backward compatible to the original version. The availability of the source code of LECCS has allowed for extending the tools by the new XPP operations in the described way.

The development of the XPP algorithms have to be done with separate tools, provided by PACT Corp.

Application Results

As a first analysis application, an inverse Discrete Cosine Transform (DCT) applied to an 8x8 pixel block was implemented. For all simulations, a 2.50 MHz clock frequency for the LEON processor and a 50 MHz clock frequency for XPP was used. The usage of XPP accelerates the computation of the iDCT by about a factor of four, depending on the communication mode. However, XPP has to be configured before computing the iDCT on it. The following table shows the configuration time for this algorithm.

	<u>LEON alone</u>	<u>LEON with XPP in IRQ Mode</u>	<u>LEON with XPP in Poll Mode</u>	<u>LEON with XPP in Hold Mode</u>
<u>Configuration of XPP</u>	————	<u>71.308 ns</u> <u>17.827 cycles</u>	<u>84.364 ns</u> <u>21.091 cycles</u>	<u>77.976 ns</u> <u>19.494 cycles</u>
<u>2D iDCT (8x8)</u>	<u>14.672 ns</u> <u>3.668 cycles</u>	<u>3.272 ns</u> <u>818 cycles</u>	<u>3.872 ns</u> <u>968 cycles</u>	<u>3.568 ns</u> <u>892 cycles</u>

As shown in Fig. 7, the benefit brought by XPP rises with the number of iDCT blocks computed by it before reconfiguration. Accordingly, the number of reconfigurations during complex algorithms should be minimized.

5

A first complex application implemented on the system is MPEG-4 decoding. The optimization of the algorithm partitioning on LEON and XPP is still in progress. Fig. 8 is a block diagram of the MPEG-4 decoding algorithm. Frames with 320 x 240 pixels were decoded. LEON, by using SPARC V8 standard instructions, decodes one frame in 23.46 seconds. In a first implementation of MPEG-4 using the XPP, only the iDCT is computed by XPP. The rest of the MPEG-4 decoding is still done with LEON. With the help of XPP, one frame is decoded in 17.98s. This is a performance boost of more than twenty percent. Since the XPP performance gain by accelerating the iDCT algorithm only is very low at the moment, we work on XPP implementations of Huffmann-decoding, dequantization, and prediction decoding. So the performance boost of this implementation against the standalone LEON will be increased.

10

15

2——HARDWARE

2.1——Design Parameter Changes

20 Since ~~For integration of the XPP core shall be integrated~~ as a functional unit into a standard RISC core, some system parameters ~~have to~~ may be reconsidered as follows:

2.1.1——Pipelining / Concurrency / Synchronicity

25 RISC instructions of totally different type (Ld/St, ALU, MuL/Div/MAC, FPALU, FPMul-, etc.) ~~are~~ may be executed in separate specialized functional units to increase the fraction of silicon that is busy on average. Such functional unit separation has led to superscalar RISC designs; that exploit higher levels of parallelism.

Each functional unit of a RISC core ~~is~~may be highly pipelined to improve throughput. Pipelining ~~overlaps~~may overlap the execution of several instructions by splitting them into unrelated phases, which ~~are~~may be executed in different stages of the pipeline. Thus, different stages of consecutive instructions can be executed in parallel with each stage taking
5 much less time to execute. This ~~allows~~may allow higher core frequencies.

~~Since~~With an approximate subdivision of the pipelines of all functional units ~~are~~approximately subdivided into sub-operations of the same size (execution time), these functional units / pipelines may execute in a highly synchronous manner with complex
10 floating point pipelines being the exception.

Since the XPP core uses ~~data flow~~data flow computation, it is pipelined by design. However, a single configuration usually implements a loop of the application, so the configuration remains active for many cycles, unlike the instructions in every other functional unit, which
15 typically execute for one or two cycles at most. Therefore, it is still worthwhile to consider the separation of several phases, (e.g., Ld / Ex / Store), of an XPP configuration, ([[=]] i.e., an XPP instruction), into several functional units to improve concurrency via pipelining on this coarser scale. This also ~~improves~~may improve throughput, and response time in conjunction with multi tasking operations and implementations of simultaneous
20 multithreading (SMT).

The multi cycle execution time may also ~~forbids~~forbid a strongly synchronous execution scheme and may rather ~~leads~~lead to an asynchronous scheme, e.g., like for e.g. floating point square root units. This in turn ~~necessitates~~may necessitate the existence of explicit
25 synchronization instructions.

2.1.2—Core ~~frequency~~Frequency / Memory Hierarchy

As a functional unit, the XPP's operating frequency ~~will~~may either be half of the core frequency or equal to the core frequency of the RISC. Almost every RISC core currently on
30 the market exceeds its memory bus frequency with its core frequency by a larger factor. Therefore, caches are employed, forming what is commonly called the memory hierarchy: Each, where each layer of cache is larger but slower than its predecessors.

This memory hierarchy does not help to speed up computations which shuffle large amounts of data, with little or no data reuse. These computations are called “bounded by memory bandwidth.” However, other types of computations with more data locality (another ~~name~~term for data reuse) may gain performance as long as they fit into one of the upper layers of the memory hierarchy. This is the class of applications that ~~gain~~gains the highest speedups when a memory hierarchy is introduced.

Classical vectorization can be used to transform memory-bounded algorithms, with a data set too big to fit into the upper layers of the memory hierarchy. Rewriting the code to reuse smaller data sets sooner exposes memory reuse on a smaller scale. As the new data set size is chosen to fit into the caches of the memory hierarchy, the algorithm is not memory bounded ~~any more~~anymore, yielding significant speed-ups.

2.1.3—Software / Multitasking Operating Systems

As the XPP is introduced into a RISC core, the changed environment - higher frequency and the memory hierarchy - may necessitate, not only reconsideration of hardware design parameters, but also a reevaluation of the software environment.

Memory Hierarchy

The introduction of a memory hierarchy ~~enhances~~may enhance the set of applications that can be implemented efficiently. So far, the XPP has mostly been used for algorithms that read their data sets in a linear manner, applying some calculations in a pipelined fashion and writing the data back to memory. As long as all of the computation fits into the XPP array, these algorithms are memory bounded. Typical applications are filtering and audio signal processing in general.

But there is another set of algorithms; that have even higher computational complexity and higher memory bandwidth requirements. -Examples are picture and video processing, where a second and third dimension of data coherence opens up. This coherence is, e.g., exploited by picture and video compression algorithms; that scan pictures in both dimensions to find similarities, even searching consecutive pictures of a video stream for analogies. ~~Naturally these~~ These algorithms have a much higher algorithmic complexity as well as higher memory requirements. Yet they are data local, either by design or ~~they can be transformed to~~

~~be~~by transformation, thus efficiently exploiting the memory hierarchy and the higher clock frequencies of processors with memory hierarchies.

Multi Tasking

5 The introduction into a standard RISC core makes it necessary to understand and support the needs of a multitasking operating system, as standard RISC processors are usually operated in multitasking environments. With multitasking, the operating system ~~switches~~may switch the executed application on a regular basis, thus simulating concurrent execution of several applications (tasks). To switch tasks, the operating system ~~has~~may have to save the state,
10 (e.g., the contents of all registers), of the running task and then reload the state of another task. Hence, it ~~is~~may be necessary to determine what the state of the processor is, and to keep it as small as possible to allow efficient context switches.

~~Modern~~Modern microprocessors gain their performance from multiple specialized and deeply
15 pipelined functional units and high memory hierarchies, enabling high core frequencies. But high memory hierarchies mean that there is a high penalty for cache misses- due to the difference between core and memory frequency. Many core cycles may pass until the values are finally available from memory. Deep pipelines incur pipeline stalls due to data dependencies as well as branch penalties for mispredicted conditional branches. Specialized
20 functional units like floating point units idle for integer-only programs. For these reasons, average functional unit utilization is much too low.

The newest development with RISC processors, Simultaneous MultiThreading (SMT), adds hardware support for a finer granularity (instruction / functional unit level) switching of tasks,
25 exposing more than one independent instruction stream to be executed. Thus, whenever one instruction stream stalls or doesn't utilize all functional units, the other one can jump in. This improves functional unit utilization for today's processors.

With SMT, the task (process) switching is done in hardware, so the processor state has to be
30 duplicated in hardware. So again it is most efficient to keep the state as small as possible. For the combination of the PACT XPP and a standard RISC processor, SMT ~~is~~may be very beneficial, since the XPP configurations may execute longer than the average RISC instruction. Thus, another task can utilize the other functional units, while a configuration is

running. On the other ~~side~~hand, not every task will utilize the XPP, so while one such non-XPP task is running, another one will be able to use the XPP core.

~~2.2~~ Communication Between the RISC Core and the XPP Core

- 5 ~~In the~~The following section ~~introduces~~are several possible embodiments that are each a possible hardware implementationsimplementation for accessing memory.

~~2.2.1~~ Streaming

- 10 Since streaming can only support (number_of_IO_ports * width_of_IO_port) bits per cycle, it ~~is only~~may be well suited for only small XPP arrays with heavily pipelined configurations that feature few inputs and outputs. As the pipelines take a long time to fill and empty while the running ~~time~~time of a configuration is limited (as described ~~under~~herein with respect to “context switches”), this type of communication does not scale well to bigger XPP arrays and XPP frequencies near the RISC core frequency.

15

■ **Streaming from the RISC core**

- In this setup, the RISC ~~supplies~~may supply the XPP array with the streaming data. Since the RISC core ~~has~~may have to execute several instructions to compute addresses and load an item from memory, this setup is only suited, if the XPP core is reading data with a frequency
20 much lower than the RISC core frequency.

■ **Streaming via DMA**

- In this mode the RISC core only initializes a DMA channel which may then ~~supplies~~supply the data items to the streaming port of the XPP core.

25

~~2.2.2~~ Shared Memory (Main Memory)

- In this configuration, the XPP array configuration ~~uses~~may use a number of PAEs to generate an address that ~~is~~ used to access main memory through the IO ports. As the number of IO ports ~~is~~may be very limited, this approach ~~suffers~~may suffer from the ~~same~~same limitations as
30 the previous one, although for larger XPP arrays there is less impact of using PAEs for address generation. However, this approach ~~is~~may still be useful for loading values from very ~~sparse~~sparse vectors.

~~2.2.3~~ Shared Memory (IRAM)

This data access mechanism uses the IRAM elements to store data for local computations. The IRAMs can either be viewed as vector registers or as local copies of main memory.

~~There~~The following are several ways in which to fill the IRAMs with data:

- 5 1. The IRAMs ~~are~~may be loaded in advance by a separate configuration using streaming.

This method can be implemented with the current XPP architecture. The IRAMs act as vector registers. As explicated above, this ~~will~~may limit the performance of the XPP array, especially as the IRAMs will always be part of the externally visible state and hence must be saved and ~~re-stored~~restored on context switches.

10

2. The IRAMs ~~can~~may be loaded in advance by separate load-instructions. This is similar to the first method. Load-instructions ~~are~~may be implemented in hardware which ~~load~~loads the data into the IRAMs. The load-instructions can be viewed as a hard coded load configuration. Therefore, configuration reloads ~~are~~may be reduced. Additionally, the special load- instructions may use a wider interface to the memory hierarchy. Therefore, a more efficient method than streaming can be used.
- 15

- 20 3. The IRAMs can be loaded by a “burst preload from memory” instruction of the cache controller. No configuration or load-instruction is needed on the XPP. The IRAM load ~~is~~may be implemented in the cache controller and triggered by the RISC processor. But the IRAMs may still act as vector registers and ~~are~~may be therefore included in the externally visible state.

- 25 4. The best mode, however ~~is~~, may be a combination of the previous solutions with the extension of a cache:
A preload instruction ~~maps~~may map a specific memory area defined by starting address and size to an IRAM. This ~~triggers~~may trigger a (delayed, low priority) burst load from the memory hierarchy (cache):- After all IRAMs are mapped, the next configuration can be activated. The activation ~~incurs~~may incur a wait until all burst loads are completed. However, if the preload instructions are issued long enough in advance and no interrupt or task switch destroys cache locality, the wait will not consume any time.
- 30

To specify a memory block as output-only IRAM, a “preload clean” instruction ~~is~~may be used, which ~~avoids~~may avoid loading data from memory. The “preload clean” instruction just indicates the IRAM for write back.

5 A synchronization instruction ~~is~~may be needed to make sure that the content of a specific memory area, which is cached in IRAM, is written back to the memory hierarchy. This can be done globally (full write back), or selectively by specifying the memory area, which will be accessed.

~~2.3~~—State of the XPP Core

10 As ~~described in the previous section~~discussed above, the size of the state ~~is~~may be crucial for the efficiency of context switches. However, although the size of the state ~~is~~may be fixed for the XPP core, ~~it depends whether or not they have to be saved~~ may depend on the declaration of the various state elements, ~~whether they have to be saved or not~~.

15 The state of the XPP core can be classified as:

~~1~~1. Read only (instruction data)

- ~~■~~——configuration data, consisting of PAE configuration and routing configuration data; and

~~2~~2. Read - Write

- 20
- ~~■~~——the contents of the data registers and latches of the PAEs, which are driven onto the busses
 - ~~■~~——the contents of the IRAM elements.

~~2.3.1~~—Limiting Memory Traffic

25 There are several possibilities to limit the amount of memory traffic during context switches, as follows:

~~Do not save read-only data~~

Do Not Save Read-Only Data

30 This ~~avoids~~may avoid storing configuration data, since configuration data is read only. The current configuration ~~is~~may be simply overwritten by the new one.

~~Save less data~~Less Data

If a configuration is defined to be uninterruptible (non pre-emptive), all of the local state on the busses and in the PAEs can be declared as scratch. This means that every configuration ~~gets~~may get its input data from the IRAMs and ~~writes~~may write its output data to the IRAMs. So after the configuration has finished, all information in the PAEs and on the buses ~~is~~may be redundant or invalid and ~~does~~saving of the information might not ~~have to be saved~~be required.

Save ~~modified data only~~ Modified Data Only

To reduce the amount of R/W data, which has to be saved, ~~we need to~~the method may keep track of the modification state of the different entities. This ~~incurs~~may incur a silicon area penalty for the additional “dirty” bits.

Use ~~eaching~~ Caching to reduce Reduce the memory traffic Memory Traffic

The configuration manager ~~handles~~may handle manual preloading of configurations.

Preloading ~~will~~may help in parallelizing the memory transfers with other computations

during the task switch. This cache can also reduce the memory traffic for frequent context switches, provided that a Least Recently Used (LRU) replacement strategy is implemented in addition to the preload mechanism.

The IRAMs can be defined to be local cache copies of main memory as ~~proposed as fourth~~

~~method in section 2.23~~discussed above under the heading “Shared Memory (IRAM).” Then

each IRAM ~~is~~may be associated with a starting address and modification state information.

The IRAM memory cells ~~are~~may be replicated. An IRAM PEA ~~contains~~PAE may contain an

IRAM block with multiple IRAM instances. ~~Only~~It may be that only the starting addresses

of the IRAMs have to be saved and restored as context. The starting addresses for the

IRAMs of the current configuration select the IRAM instances with identical addresses to be used.

If no address tag of an IRAM instance matches the address of the newly, loaded context, the corresponding memory area ~~is~~may be loaded to an empty IRAM instance.

If no empty IRAM instance is available, a clean (unmodified) instance ~~is~~may be declared empty (and hence ~~must~~it may be required for it to be reloaded later on).

If no clean IRAM instance is available, a modified (dirty) instance ~~is~~may be cleaned by writing its data back to main memory. This ~~adds~~may add a certain delay for the write back.

5 This delay can be avoided, if a separate state machine (cache controller) tries to clean inactive IRAM instances by using unused memory cycles to write back the IRAM instances' contents.

2.4—Context Switches

Usually a processor is viewed as executing a single stream of instructions. But today's multi-tasking operating systems support hundreds of tasks being executed on a single processor.
10 This is achieved by switching contexts, where all, or at least the most relevant parts, of the processor state, which belong to the current task - the task's context - is exchanged with the state of another task, that will be executed next.

There are three types of context switches: switching of virtual processors with simultaneous
15 ~~multi-threading~~multithreading (SMT, also known as HyperThreading), execution of an Interrupt Service Routine (ISR), and a Task Switch.

2.4.1—SMT Virtual Processor Switch

This type of context switch ~~is~~may be executed without software interaction, totally in
20 hardware. Instructions of several instruction streams are merged into a single instruction stream to increase instruction level parallelism and improve functional unit utilization. Hence, the processor state cannot be stored to and reloaded from memory between instructions from different instruction streams: ~~Imagine the worst case.~~ For example, in an instance of alternating instructions from two streams and the hundreds to thousandthousands
25 of cycles might be needed to write the processor state to memory and read in another state.

Hence hardware designers have to replicate the internal state for every virtual processor. Every instruction ~~is~~may be executed within the context (on the state) of the virtual processor, whose program counter was used to fetch the instruction. By replicating the state, only the
30 multiplexers, which have to ~~b~~be inserted to select one of the different states, have to be switched.

Thus the size of the state ~~may also increases~~increase the silicon area needed to implement SMT, so the size of the state ~~is~~may be crucial for many design decisions.

2.4.2—Interrupt Service Routine

This type of context switch ~~is~~may be handled partially by hardware and partially by software.

~~All~~ It may be required for all of the state modified by the ISR ~~has-to be saved on entry and~~

5 ~~must~~it may be required for it to be restored on exit.

The part of the state, which is destroyed by the jump to the ISR, ~~is~~ may be saved by hardware, (e.g., the program counter). It ~~is~~may be the ISR's responsibility to save and restore the state of all other resources, that are actually used within the ISR.

10

The more state information to be saved, the slower the interrupt response time ~~will~~may be and the greater the performance impact ~~will~~may be if external events trigger interrupts at a high rate.

15 The execution model of the instructions ~~will~~may also affect the tradeoff between short interrupt latencies and maximum throughput. Throughput ~~is~~may be maximized if the instructions in the pipeline are finished, and the instructions of the ISR are chained. This may adversely ~~affects~~affect the interrupt latency. If, however, the instructions are abandoned (pre-empted) in favor of a short interrupt latency, ~~they must~~it may be required for them to be

20 fetched again later, which ~~affects~~may affect throughput. The third possibility would be to save the internal state of the instructions within the pipeline, but this ~~requires~~may require too much hardware effort. Usually this is not done.

2.4.3—Task Switch

25 This type of context switch ~~is~~may be executed totally in software. ~~All~~ It may be required for all of a task's context (state) ~~has-to be saved to memory, and it may be required for the~~ context of the new task ~~has-to be reloaded~~. Since tasks are usually allowed to use all of the processor's resources to achieve top performance, it may be required to save and restore all of the processor state ~~has-to be saved and restored~~. If the amount of state is excessive, it may be

30 required for the rate of context switches ~~must~~to be decreased by less frequent rescheduling, or a severe throughput degradation ~~will~~may result, as most of the time ~~will~~may be spent in saving and restoring task contexts. This in turn ~~increases~~may increase the response time for the tasks.

2.5—A Load Store Architecture

- ~~We propose~~In an example embodiment of the present invention, an XPP integration may be provided as an asynchronously pipelined functional unit for the RISC. ~~We further propose an~~ An explicitly preloaded cache may be provided for the IRAMs, on top of the memory hierarchy existing within the RISC (as ~~proposed as fourth method in section 2.2.3~~).discussed above under the heading “Shared Memory (IRAM).” Additionally a de-centralized explicitly preloaded configuration cache within the PAE array is~~may be~~ employed to support preloading of configurations and fast switching between configurations.
- 10 Since the IRAM content is an explicitly preloaded memory area, a virtually unlimited number of such IRAMs can be used. They ~~are~~may be identified by their memory address and their size. The IRAM content ~~is~~may be explicitly preloaded by the application. Caching ~~will~~may increase performance by reusing data from the memory hierarchy. The cached operation may also ~~eliminate~~eliminate the need for explicit store instructions; they ~~are~~may be handled
- 15 implicitly by cache write back operations but can also be forced to synchronize with the RISC.

- The pipeline stages of the XPP functional unit ~~are~~may be Load, Execute, and Write Back (Store). The store ~~is~~may be executed delayed as a cache write back. The pipeline stages may
- 20 execute in an asynchronous fashion, thus hiding the variable delays from the cache preloads and the PAE array.

- The XPP functional unit ~~is~~may be decoupled of the RISC by a FIFO, ~~which is~~ fed with the XPP instructions. At the head of this FIFO, the XPP PAE ~~consumes~~may consume and
- 25 ~~executes~~execute the configurations and the preloaded IRAMs. Synchronization of the XPP and the RISC is~~may be~~ done explicitly by a synchronization instruction.

Instructions

- ~~In the following we define the~~Embodiments of the present invention may require certain
- 30 instruction formats needed for the proposed architecture. We use. Data types may be specified using a C style prototype definition ~~to specify data types. All instructions, except the XPPSync instruction.~~ The following are example instruction formats which may be required, all of which execute asynchronously. ~~The, except for an XPPSync instruction, which~~ can be used to force synchronization.

XPPPreloadConfig (void *ConfigurationStartAddress)

The configuration is~~may be~~ added to the preload FIFO to be loaded into the configuration cache within the PAE array.

- 5 Note that speculative preloads ~~are~~is possible; since successive preload commands overwrite the previous.

The parameter is a pointer register of the RISC pointer register file. The size is implicitly contained in the configuration.

10

XPPPreload (int IRAM, void *StartAddress, int Size)

XPPPreloadClean (int IRAM, void *StartAddress, int Size)

This instruction ~~specifies~~may specify the contents of the IRAM ~~for~~for the next configuration execution. In fact, the memory area is~~may be~~ added to the preload FIFO to be loaded into the specified IRAM.

15

The first parameter is~~may be~~ the IRAM number. This is~~may be~~ an immediate (constant) value.

- 20 The second parameter is~~may be~~ a pointer to the starting address. This parameter is~~may be~~ provided in a pointer register of the RISC pointer register file.

The third parameter is~~may be~~ the size in units of 32 bit words. This is~~may be~~ an integer value. It ~~resides~~may reside in a general- purpose register of the RISC's integer register file.

25

The first variant may actually ~~preloads~~preload the data from memory.

The second variant is~~may be~~ for write-only accesses. It ~~skips~~may skip the loading operation.

Thus, it may be that no cache misses can occur for this IRAM. Only the address and size are

- 30 defined. They are obviously needed for the write back operation of the IRAM cache.

Note that speculative preloads are possible; since successive preload commands to the same IRAM overwrite each other (if no configuration is executed ~~inbetween~~in between). Thus,

only the last preload command is~~may be~~ actually effective; when the configuration is executed.

XPPEXecute ()

- 5 This instruction ~~executes~~may execute the last preloaded configuration with the last preloaded IRAM contents. Actually, a configuration start command is~~may be~~ issued to the FIFO. Then the FIFO is~~may be~~ advanced; ~~this means~~. This may mean that further preload commands will specify the next configuration or parameters for the next configuration.
- 10 Whenever a configuration finishes, the next one is~~may be~~ consumed from the head of the FIFO, if its start command has already been issued.

XPPSync (void *StartAddress, int Size)

- 15 This instruction ~~forces~~may force write back operations for all IRAMs that overlap the given memory area. If overlapping IRAMs are still in use by a configuration or preloaded to be used, this operation will block. Giving an address of NULL (zero) and a size of MAX_INT (bigger than the actual memory), this instruction can also be used to wait until all issued configurations finish.

20 **2.5.1—A Basic Implementation**

As shown in Fig. 9, the XPP core 102 may share a memory hierarchy with the RISC core 112 using a special cache controller 125-130.

- 25 Fig. 10 shows an IRAM and configuration cache controller data structures and a usage example (instructions).

- The preload-FIFOs in ~~the above figure~~Fig. 10 may contain the addresses and sizes for already issued IRAM ~~pre-loads~~preloads, exposing them to the XPP cache controller. The FIFOs may have to be duplicated for every virtual processor in an SMT environment. “Tag” is the
- 30 typical tag for a cache line containing starting address, size, and state (*empty / clean / dirty / in-use*). The additional *in-use* state signals usage by the current configuration. The cache controller cannot manipulate these IRAM instances.

The execute configuration command ~~advances~~may advance all preload FIFOs, copying the old state to the newly created entry. This way the following preloads may replace the previously used IRAMs and configurations. If no preload is issued for an IRAM before the configuration is executed, the preload of the previous configuration ~~is~~may be retained.

5 Therefore, it may be that it is not necessary to repeat identical preloads for an IRAM in consecutive configurations.

~~Figure 3: Asynchronous pipeline of the XPP~~

Each configuration's execute command ~~has~~may have to be delayed (stalled) until all
10 necessary preloads are finished, either explicitly by the use of a synchronization command or implicitly by the cache controller. Hence the cache controller (XPP Ld/St unit) ~~has~~125 may have to handle the synchronization and execute commands as well, actually starting the configuration as soon as all data is ready. After the termination of the configuration, dirty IRAMs ~~are~~may be written back to memory as soon as possible, if their content is not reused
15 in the same IRAM. Therefore the XPP PAE array (XPP core 102) and the XPP cache controller 125 can be seen as a single unit since they do not have different instruction streams: ~~rather~~Rather, the cache controller can be seen as the configuration fetch (CF), operand fetch (OF) (IRAM preload) and write back (WB) stage of the XPP pipeline, also triggering the execute stage (EX) (PAE array). Fig. 11 shows the asynchronous pipeline of
20 the XPP 100.

Due to the long latencies, and their non-predictability (cache misses, variable length configurations), the stages can be overlapped several configurations wide using the configuration and data preload FIFO, ([[=]] i.e., pipeline), for loose coupling. ~~So if~~ If a
25 configuration is executing and the data for the next has already been preloaded, the data for the next but one configuration ~~is~~may be preloaded. ~~These preloads can be speculative; the~~
The amount of speculation ~~is~~may be the compiler's trade-off. The reasonable length of the preload FIFO can- be several configurations; ~~it is~~. It may be limited by diminishing returns, algorithm properties, the compiler's ability to schedule preloads early and by silicon usage
30 due to the IRAM duplication factor, which ~~has~~may have to be at least as big as the FIFO length. Due to this loosely coupled operation, the interlocking ~~-(to avoid data hazards between IRAMs)~~ cannot be done optimally by software (scheduling), but ~~has~~may have to be

enforced by hardware (hardware interlocking). Hence the XPP cache controller and the XPP PAE array can be seen as separate but not totally independent functional units.

~~Figure 4: State transition diagram for the XPP cache controller~~

5 The XPP cache controller ~~has~~may have several tasks. These are depicted as states in ~~the above diagram. Fig. 12.~~ State transitions may take place along the edges between states, whenever the condition for the edge is true. As soon as the condition is not true any more, the reverse state transition ~~takes~~may take place. The activities for the states ~~are~~may be as follows:

10 At the lowest priority, the XPP cache controller ~~has~~may have to fulfill already issued preload commands, while writing back dirty IRAMs as soon as possible.

As soon as a configuration finishes, the next configuration can be started. This is a more urgent task than write backs or future preloads. To be able to do that, all associated yet unsatisfied preloads may have to be finished first. Thus, they ~~are~~may be preloaded with the
15 high priority inherited from the execute state.

A preload in turn can be blocked by an overlapping *in-use* or *dirty* IRAM instance in a different block or by the lack of *empty* IRAM instances in the target IRAM block. The former can be resolved by waiting for the configuration to finish and / or by a write back. To resolve the latter, the least recently used *clean* IRAM can be discarded, thus becoming *empty*.

20 If no *empty* or *clean* IRAM instance exists, a *dirty* one ~~has~~may have to be written back⁷ to the memory hierarchy. It cannot occur that no *empty*, *clean*, or *dirty* IRAM instances exist, since only one instance can be *in-use* and there should be more than one instance in an IRAM block; [[-]] otherwise, no caching effect is achieved.

~~Figure 5: Adding Simultaneous multithreading~~

25

In an SMT environment the load FIFOs may have to be replicated for every virtual processor. The pipelines of the functional units ~~are~~may be fed from the shared fetch / reorder / issue stage. All functional units may execute in parallel. Different units can execute instructions of different virtual processors. Fig. 13 shows adding of simultaneous multithreading.

30

So ~~we get~~ the following design parameters, with their smallest initial value, may be obtained:

- IRAM length:

128 words

The longer the IRAM length, the longer the running time of the configuration and the less influence the pipeline startup has.

- FIFO length: 1
This parameter ~~helps~~may help to hide cache misses during preloading. The longer
5 the FIFO length, the less disruptive is a series of cache misses for a single
configuration.
- IRAM duplication factor: (pipeline stages+caching factor)*virtual processors: 3
Pipeline stages is the number of pipeline stages LD/EX/WB plus one for every
FIFO stage above one: 3
10 Caching factor is the number of IRAM duplicates available for caching: 0
Virtual processors is the number of virtual processors with SMT: 1

The size of the state of a virtual processor is mainly dependent on the FIFO length. It is:—
FIFO length * #IRAM ports * (32 bit (Address) + 32 bit (Size)).

15 This ~~has~~may have to be replicated for every virtual processor.

The total size of memory used for the IRAMs ~~is~~may be:

#IRAM ports * ~~IRAM~~IRAM duplication factor* IRAM length * 32 bit.

20 A first implementation will probably keep close to the above-stated minimum parameters, using a FIFO length of one, an IRAM duplication factor of four, an IRAM length of 128 and no simultaneous multithreading.

25 **2.5.2—Implementation Improvements**

Write Pointer

To further decrease the penalty for unloaded IRAMs, a simple write, pointer may be used per IRAM, which ~~keeps~~may keep track of the last address already in the IRAM. Thus, no stall is required, unless an access beyond this write pointer is encountered. This ~~is~~may be especially
30 useful, if all IRAMs have to be reloaded after a task switch. The delay to the configuration start can be much shorter, especially, if the preload engine of the cache controller, chooses the blocking IRAM next whenever several IRAMs need further loading.

Longer FIFOs

The frequency at the bottom of the memory hierarchy (main memory) cannot be raised to the same extent as the frequency of the CPU core. To increase the concurrency between the RISC core 112 and the PACT XPP core, 102, the prefetch FIFOs in ~~the above drawing~~ Fig. 13

5 can be extended. Thus, the IRAM contents for several configurations can be preloaded, like the configurations themselves. A simple convention makes clear which IRAM preloads belong to which configuration: ~~the~~ The configuration execute switches to the next configuration context. This can be accomplished by advancing the FIFO write pointer with every configuration execute, while leaving it unchanged after every preload: Unassigned
10 IRAM FIFO entries may keep their contents from the previous configuration, so every succeeding configuration ~~will~~ may use the preceding configuration's IRAMx if no different IRAMx was preloaded.

If none of the memory areas to be copied to IRAMs is in any cache, extending the FIFOs
15 does not help, as the memory *is* the bottleneck. So the cache size should be adjusted together with the FIFO length.

A drawback of extending the FIFO length is the increased likelihood that the IRAM content written by an earlier configuration is reused by a later one in another IRAM. A cache
20 coherence protocol can clear the situation. Note, however, that the situation can be resolved more easily: If an overlap between any new IRAM area and a currently dirty IRAM contents of another IRAM bank is detected, the new IRAM is simply not loaded until the write back of the changed IRAM has finished. Thus, the execution of the new configuration
is may be delayed until the correct data is available.

25 For a short (single entry) FIFO, an overlap is extremely unlikely, since the compiler will usually leave the output IRAM contents of the previous configuration in place for the next configuration to skip the preload. The compiler ~~does~~ may do so using a coalescing algorithm for the IRAMs / vector registers. The coalescing algorithm ~~is~~ may be the same as used for
30 register coalescing in register allocation.

Read Only ~~IRAMs~~ IRAMS

Whenever the memory, that is used by the executing configuration, is the source of a preload command for another IRAM, an XPP pipeline stall ~~occurs~~ may occur. The preload can only

be started, when the configuration has finished, and—, if the content was modified—, the memory content has been written to the cache. To decrease the number of pipeline stalls, it ~~is~~may be beneficial to add an additional read- only IRAM state. If the IRAM is read only, the content cannot be changed, and the preload of the data to the other IRAM can proceed without delay. This ~~requires~~may require an extension to the preload instructions. The XppPreload and the XppPreloadClean instruction formats can be combined to a single instruction format, that has two additional bits, stating whether the IRAM will be read and/or written. To support debugging, violations should be checked at the IRAM ports, raising an exception when needed.

2.5.3—Support for Data Distribution and Data Reorganization

The IRAMs ~~are~~may be block-oriented structures, which can be read in any order by the PAE array. However, the address generation, ~~adds~~ may add complexity, reducing the number of PAEs available for the actual computation. ~~So it is best, if~~ Accordingly, the IRAMs ~~are~~may be accessed in linear order. The memory hierarchy ~~is~~may be block oriented as well, further encouraging linear access patterns in the code to avoid cache misses.

As the IRAM read ports limit the bandwidth between each IRAM and the PAE array to one word read per cycle, it can be beneficial to distribute the data over several IRAMs to remove this bottleneck. The top of the memory hierarchy is the source of the data, so the ~~amount~~number of cache misses never increases when the access pattern is changed, as long as the data locality is not destroyed.

Many algorithms access memory in linear order by definition to utilize block reading and simple address calculations. In most other cases and in the cases where loop tiling is needed to increase the data bandwidth between the IRAMs and the PAE array, the code can be transformed in a way that data is accessed in optimal order. In many of the remaining cases, the compiler ~~can~~can modify the access pattern by data layout rearrangements, (e.g., array merging), so that finally the data is accessed in the desired pattern. If none of these optimizations can be used because of dependencies, or because the data layout is fixed, there are still two possibilities to improve performance:—, which are data duplication and data reordering.

Data Duplication

Data ~~is~~may be duplicated in several IRAMs. This ~~circumvents~~may circumvent the IRAM read port bottleneck, allowing several data items to be read from the input every cycle.

5 Several options are possible with a common drawback: ~~data~~. Data duplication can only be applied to input data: ~~output~~. Output IRAMs obviously cannot have overlapping address ranges.

• ~~o~~—Using several IRAM preload commands specifying just different target IRAMs:
10 This way cache misses may occur only for the first preload. All other preloads ~~will~~may take place without cache misses—~~only~~. Only the time to transfer the data from the top of the memory hierarchy to the IRAMs is needed for every additional load. This is only beneficial; if the cache misses plus the additional transfer times do not exceed the execution time for the configuration.

15

• ~~o~~—Using an IRAM preload instruction to load multiple IRAMs concurrently:
As identical data is needed in several IRAMs, they can be loaded concurrently by writing the same values to all of them. This amounts to finding a clean IRAM instance for every target IRAM, connecting them all to the bus, and writing the data to
20 the bus. The problem with this instruction ~~is~~may be that it requires a bigger immediate field for the destination (16 bits instead of 4 for the XPP 64). Accordingly, this instruction format ~~grows~~may grow at a higher rate; when the number of IRAMs is increased for bigger XPP arrays.

25 The interface of this instruction ~~looks like~~is for example:

XPPPreloadMultiple (int IRAMS, void *StartAddress, int Size).

This instruction ~~behaves~~may behave as the XPPPreload / XPPPreloadClean instructions with the exception of the first parameter: The first parameter is IRAMS. This ~~is~~may be an
30 immediate (constant) value. The value ~~is~~may be a bitmap—~~for~~. For every bit in the bitmap, the IRAM with that number ~~is~~may be a target for the load operation.

There is no “clean” version, since data duplication is applicable for read data only.

Data Reordering

Data reordering changes the access pattern to the data, only. It does not change the amount of memory that is read. Thus, the number of cache misses ~~stays~~may stay the same.

- ~~o~~ Adding additional functionality to the hardware:

5 ~~o~~ Adding a vector stride to the preload instruction.

A *stride* (displacement between two elements in memory) ~~is~~may be used in vector load operations to load, e.g., a column of a matrix into a vector register.

10 This is still a linear access pattern. It can be implemented in hardware by giving a stride to the preload instruction and adding the stride to the IRAM identification state. One problem with this instruction ~~is~~may be that the number of possible cache misses per IRAM load rises~~:-~~. In the worst case it can be one cache miss per loaded value; if the stride is equal to the cache line size and all data is not in the cache. But as already stated~~:-~~, the total number of misses stays the same—~~just~~. Just the distribution

15 changes. Still, this is an undesirable effect.

The other problem ~~is~~may be the complexity of the implementation and a possibly limited throughput, as the data paths between the layers of the memory hierarchy are optimized for block transfers. Transferring non-contiguous words will not use wide

20 busses in an optimal fashion.

The interface of the instruction ~~looks like~~is for example:

XPPPreloadStride (int IRAM, void *StartAddress, int Size, int Stride)

25 XPPPreloadCleanStride (int IRAM, void *StartAddress, int Size, int Stride).

This instruction ~~behaves~~may behave as the XPPPreload / XPPPreloadClean instructions with the addition of another parameter~~:-~~. The fourth parameter is the vector stride. This ~~is~~may be an immediate (constant) value. It ~~tells~~may tell the cache

30 controller, to load only every n^{th} value to the specified IRAM.

- ~~o~~ Reordering the data at run time, introducing temporary copies.

~~o~~ On the RISC:

The RISC can copy data at a maximum rate of one word per cycle for simple address computations and at a somewhat lower rate for more complex ones.

5 With a memory hierarchy, the sources ~~will~~may be read from memory (or cache, if they were used recently) once and written to the temporary copy, which ~~will~~may then reside in the cache, too. This ~~increases~~may increase the pressure in the memory hierarchy by the amount of memory used for the temporaries. Since temporaries are allocated on the stack memory, which ~~is~~may be re-used frequently, the chances are good that the dirty memory area is ~~re-defined~~redefined before it is written back to
10 memory. Hence the write back operation to memory is of no concern.

• ~~o~~ Via an XPP configuration:

15 The PAE array can read and write one value from every IRAM per cycle. Thus, if half of the IRAMs are used as inputs and half of the IRAMs are used as outputs, up to eight (or more, depending on the number of IRAMs), values can be reordered per cycle, using the PAE array for address generation. As the inputs and outputs reside in IRAMs, it does not matter, if the reordering is done before or after the configuration—~~that uses the data—the~~. The IRAMs can be reused immediately.

20 IRAM Chaining

If the PAEs do not allow further unrolling, but there are still IRAMs left unused, it ~~is~~may be possible to load additional blocks of data into these IRAMs and chain two IRAMs ~~by means of~~via an address selector. This ~~does~~might not increase throughput as much as unrolling would do, but it still ~~helps~~may help to hide long pipeline startup delays whenever unrolling is
25 not possible.

2.6—Software / Hardware Interface

30 According to the design parameter changes and the corresponding changes to the hardware, ~~the hardware~~according to embodiments of the present invention, the hardware / software interface has changed. In the following ~~the most, some~~ prominent changes and their handling ~~will be~~are discussed:

2.6.1—Explicit Cache

The proposed cache is not a usual cache, which would be ~~not~~, without considering performance issues ~~in-visible, invisible~~ to the programmer / compiler, as its operation is transparent. The proposed cache is an explicit cache. Its state ~~has~~may have to be maintained
5 by software.

Cache Consistency and Pipelining of Preload / Configuration / Write back

The software ~~is~~may be responsible for cache consistency. It ~~is~~may be possible to have several IRAMs caching the same, or overlapping memory areas. As long as only one of the
10 IRAMs is written, this is perfectly ok~~:-~~. Only this IRAM will be dirty and will be written
back to memory. If, however, more than one of the IRAMs is written, ~~it~~which data will be
written to memory is not defined, ~~which data will be written to memory~~. This is a software
bug (non -deterministic behavior).

15 As the execution of the configuration is overlapped with the preloads and write backs of the
IRAMs, it ~~is~~may be possible to create preload / configuration sequences, that contain data
hazards. As the cache controller and the XPP array can be seen as separate functional units,
which are effectively pipelined, these data hazards are equivalent to pipeline hazards of a
normal instruction pipeline. As with any ordinary pipeline, there are two possibilities to
20 resolve this~~:-~~, which are hardware interlocking and software interlocking.

- Hardware interlocking:

Interlocking ~~is~~may be done by the cache controller~~:-~~. If the cache controller detects, that
the tag of a dirty or in-use item in IRAMx overlaps a memory area used for another
25 IRAM preload, it ~~has~~may have to stall that preload, effectively serializing the execution
of the current configuration and the preload.

- Software interlocking:

If the cache controller does not enforce interlocking, the code generator ~~has~~may have to
30 insert explicit synchronize instructions to take care of potential interlocks. Inter-
procedural and inter -modular alias- and data- dependency analyses can determine if this
is the case, while scheduling algorithms may help to alleviate the impact of the necessary
synchronization instructions.

In either case, as well as in the case of pipeline stalls due to cache misses, SMT can use the computation power; that would be wasted otherwise.

Code Generation for the Explicit Cache

5 Apart from the explicit synchronization instructions issued with software interlocking, the following instructions may have to be issued by the compiler.

- Configuration preload instructions, preceding the IRAM preload instructions, that will be used by that configuration. These should be scheduled as early as possible by the
10 instruction scheduler.
- IRAM preload instructions, which should also be scheduled as early as possible by the instruction scheduler.
- 15 • Configuration execute instructions, following the IRAM preload instructions for that configuration. These instructions should be scheduled between the estimated minimum and the estimated maximum of the cumulative latency of their preload instructions.
- IRAM synchronization instructions, which should be scheduled as late as possible by the
20 instruction scheduler. These instructions must be inserted before any potential access of the RISC to the data areas that are duplicated and potentially modified in the IRAMs. Typically, these instructions will follow a long chain of computations on the XPP, so they will not significantly decrease performance.

25 Asynchronicity to ~~other~~Other Functional Units

~~A~~An XppSync() must be issued by the compiler, if an instruction of another functional unit (mainly the Ld/St unit) can access a memory area; that is potentially dirty or in-use in an IRAM. This ~~forces~~may force a synchronization of the instruction streams and the cache contents, avoiding data hazards. A thorough inter-procedural and inter-modular array alias
30 analysis ~~limits~~may limit the frequency of these synchronization instructions to an acceptable level.

2.7 — Another Implementation

For the previous design, the IRAMs are existent in silicon, duplicated several times to keep the ~~pipe-line~~pipeline busy. This ~~amounts~~may amount to a large silicon area, that is not fully busy all the time, especially, when the PAE array is not used, but as well whenever the configuration does not use all of the IRAMs present in the array. The duplication may also ~~makes~~make it difficult to extend the lengths of the IRAMs, as the total size of the already large IRAM area scales linearly.

For a more silicon efficient implementation, ~~we should integrate the~~ IRAMs may be integrated into the first level cache, making this cache bigger. This means, that ~~we have to extend~~ the first level cache controller is extended to feed all IRAM ports of the PAE array. This way, the XPP and the RISC ~~will~~may share the first level cache in a more efficient manner. Whenever the XPP is executing, it ~~will~~may steal as much cache space as it needs from the RISC. Whenever the RISC alone is running it will have plenty of additional cache space to improve performance.

The PAE array ~~has~~may have the ability to read one word and write one word to each IRAM port every cycle. This can be limited to either a read or a write access per cycle, without limiting programmability. ~~If data has to be written to the same area in the same cycle,~~ another IRAM port can be used. This ~~in-creases~~may increase the number of used IRAM ports, but only under rare circumstances.

This leaves sixteen data accesses per PAE cycle in the worst case. Due to the worst case of all sixteen memory areas for the sixteen IRAM ports mapping to the same associative bank, the minimum associativity for the cache ~~is~~may be a 16-way set associativity. This ~~avoids~~may avoid cache replacement for this rare, but possible, worst-case example.

Two factors may help to support sixteen accesses per PAE array cycle:

- ~~The~~ clock frequency of the PAE array generally ~~has~~ to be lower than for the RISC by a factor of two to four. The reasons lie in the configurable routing channels with switch matrices which cannot support as high a frequency as solid point-to-point ~~aluminium~~aluminum or copper traces.

This means that two to four IRAM port accesses can- be handled serially by a single cache port, as long as all reads are serviced before all writes, if there is a potential overlap. This can be accomplished by assuming a potential overlap and enforcing a priority ordering of all accesses, giving the read accesses higher priority.

5

- A factor of two, four, or eight is possible by accessing the cache as two, four, or eight banks of lower associativity cache.

10

For a cycle divisor of four, four banks of four-way associativity will be optimal. During four successive cycles, four different accesses can be served by each bank of four way associativity. Up to four-way data duplication can be handled by using adjacent IRAM ports that are connected to the same bus (bank). For further data duplication, the data ~~has~~may have to be duplicated explicitly, using an XppPreloadMultiple() cache controller instruction. The maximum data duplication for sixteen read accesses to the same memory area is supported by an actual data duplication factor of four: one copy in each bank. This does not affect the RAM efficiency as adversely as an actual data duplication of 16 for the ~~design proposed in section 2.5.embodiment~~ discussed above under the heading “A Load Store Architecture.”

15

20

Fig. 14 shows an example of a cache structure according to an example embodiment of the present invention. The cache controller ~~is running~~may run at the same speed as the RISC. The XPP ~~is running~~may run at a lower, (e.g., quarter), speed. ~~This way~~ Accordingly, in the worst case of, sixteen read requests from the PAE array ~~need to~~may be serviced in four cycles of the cache controller, with an additional four read requests from the RISC: ~~So,~~

25

Accordingly, one bus at full speed can be used to service four IRAM read ports. Using four-way associativity, four accesses per cycle can be serviced, even in the case that all four accesses go to addresses that map to the same associative block.

30

- a) The RISC still has a 16-way set associative view of the cache, accessing all four four-way set associative banks in parallel. Due to data duplication, it is possible, that several banks return a hit. This ~~has to~~may be taken care of with a priority encoder, enabling only one bank onto the data bus:.

b) The RISC is blocked from the banks that service IRAM port accesses. Wait states are inserted accordingly.

5 c) The RISC shares the second cache access port of a two-port cache with the RAM interface, using the cycles between the RAM transfers for its accesses.

d) The cache is extended by a fifth 4-way set associative bank, used exclusively by the RISC. (The other banks are only accessed, when they are not used by the current XPP configuration. PROBLEM: dirty line in a blocked bank).

10

→ With respect to a 2 port RAM, concurrent reads ok, concurrent may be accommodated.
Concurrent R/W to a same cache line may be avoided by SW software synchronization / HW
arbitrer hardware arbiter.

15

~~Another~~ A problem is that a read could potentially address the same memory location as a write; ~~the~~. The value read ~~depends~~ may depend on the order of the operation; so that the order is fixed; i.e., all writes have to take place after all reads, but before the reads of the next cycle, except, if the reads and writes actually do not overlap. This can only be a problem with data duplication, when only one copy of the data is actually modified. Therefore, modifications are forbidden with data duplication.

20

2.7.1—Programming Model Changes

Data Interference

25 ~~With this design~~ According to an example embodiment of the present invention that is without dedicated ~~ms~~IRAMs, it is not possible ~~any more~~ anymore to load input data to the IRAMs and write the output data to a different IRAM, which is mapped to the same address, thus operating on the original, unaltered input data during the whole configuration.

30 As there are no dedicated IRAMs ~~any more~~ anymore, writes directly modify the cache contents, which will be read by succeeding reads. This changes the programming model significantly. Additional and more in-depth compiler analyses are ~~necessary~~ accordingly necessary.

2.7.2—Hiding Implementation Details

The actual number of bits in the destination field of the XppPreloadMultiple instruction is implementation dependent. It depends on the number of cache banks and their associativity, which are determined by the clock frequency divisor of the XPP PAE array relative to the cache frequency. However, this can be hidden by the assembler, ~~who translates~~which may translate IRAM ports to cache banks, thus reducing the number of bits from the number of IRAM ports to the number of banks. For the user, it is sufficient to know, that each cache bank services an adjacent set of IRAM ports starting at a power of two. Thus, it ~~is~~may be best to use data duplication for adjacent ports, starting with the highest power of two ~~bigger~~greater than the number of read ports to the duplicated area.

3—PROGRAM OPTIMIZATIONS

3.1—Code Analysis

~~In this section we describe the analyses that can~~Analyses may be performed on programs.
~~These analyses are then used by different optimizations. They to~~ describe the relationships between data and memory ~~locations~~location in the program. ~~More details can be found in several books [2,3,5]~~a program. These analyses may then be used by different optimizations.
More details regarding the analyses are discussed in Michael Wolfe, “High Performance Compilers for Parallel Computing” (Addison-Wesley 1996); Hans Zima & Barbara Chapman, “Supercompilers for parallel and vector computers” (Addison-Wesley 1991); and Steven Muchnick, “Advanced Compiler Design and Implementation” (Morgan Kaufmann 1997).

3.1.1—Data-Flow Analysis

Data-flow analysis examines the flow of scalar values through a program, to provide information about how the program manipulates its data. This information can be represented by dataflow equations that have the following general form for object i , that can be an instruction or a basic block, depending on the problem to solve:

$$Ex[i] = Prod[i] \vee (In[i] - Supp[i]).$$

~~It~~This means that data available at the end of the execution of object i , $Ex[i]$, are either produced by i , $Prod[i]$ or were alive at the beginning of i , $In[i]$, but were not deleted during the execution of i , $Supp[i]$.

These equations can be used to solve several problems ~~like~~, such as, e.g.,

- ~~■~~ — the problem of reaching definitions;_i
- ~~■~~ — the Def-Use and Use-Def chains, describing respectively, for a definition, all uses that can be reached from it, and, for a use, all definitions that can reach it;_i
- ~~■~~ — the available expressions at a point in the program;_i and/or
- 5 • ~~■~~ — the live variables at a point in the program, whose solutions are then used by several compilation phases, analysis, or optimizations.

~~As an~~ For example let us take the, with respect to a problem of computing the Def-Use chains of the variables of a program. ~~This, this~~ information can be used for instance by the data

10 dependence analysis for scalar variables or by the register allocation. A Def-Use chain is associated to each definition of a variable and is the set of all visible uses from this definition.

The data-flow equations presented above ~~iremay~~ be applied to the basic blocks to detect the variables that are passed from one block to another along the control- flow graph. ~~In the~~

~~figure below~~ In Fig. 15, which shows a control-flow graph of a piece of a program, two

15 definitions for variable x are produced: $S1$ in $B1$ and $S4$ in $B3$. Hence, the variable that can be found at the exit of $B1$ is $Ex(B1) = \{x(S1)\}$; and at the exit of $B4$ is $Ex(B4) = \{x(S4)\}$.

Moreover ~~we have~~, $Ex(B2) = Ex(B1)$ as no variable is defined in $B2$. Using these sets, ~~we~~

~~findit~~ is the case that the uses of x in $S2$ and $S3$ depend on the definition of x in ~~$B1$, $B1$~~ and that the use of x in $S5$ ~~depend~~ depends on the definitions of x in $B1$ and $B3$. The Def-use chains

20 associated with the definitions are then $D(S1) = \{S2, S3, S5\}$ and $D(S4) = \{S5\}$.

Figure. 7: Control-flow graph of a piece of program

3.1.2 Data Dependence Analysis

A data dependence graph represents the ~~dependences~~ dependencies existing between

25 operations writing or reading the same data. This graph ~~ismay~~ be used for optimizations like scheduling, or certain loop optimizations to test their semantic validity. The nodes of the

graph represent the instructions, and the edges represent the data ~~dependences~~ dependencies.

These ~~dependences~~ dependencies can be of three types: true (or flow) dependence when a variable is written before being read, anti-dependence when a variable is read before being

30 written, and output dependence when a variable is written twice. ~~Here is a~~ A more formal

definition {3} is provided in Hans Zima et al., *supra* and is presented below.

Definition

Let S and S' be 2 statements, then S' depends on S , noted $S \delta S'$ iff:

(1) ~~S is executed before S'~~

(2) ~~$\forall v \in \text{VAR} : v \in \text{DEF}(S) \wedge v \in \text{USE}(S') \vee v \in \text{USE}(S) \wedge v \in \text{DEF}(S') \vee v \in \text{DEF}(S) \wedge v \in \text{DEF}(S')$~~

5 Let S and S' be two statements. Then S' depends on S , noted $S \delta S'$ iff:

(3) ~~There is no statement T such that S is executed before T and T is executed before S' , and $v \in \text{DEF}(T)$~~

(2) $\exists v \in \text{VAR} : v \in \text{DEF}(S) \wedge v \in \text{USE}(S') \vee v \in \text{USE}(S) \wedge v \in \text{DEF}(S') \vee v \in \text{DEF}(S) \wedge v \in \text{DEF}(S')$

(3) There is no statement T such that S is executed before T and T is executed before S' ,

10 and $v \in \text{DEF}(T)$,

where VAR is the set of the variables of the program, $\text{DEF}(S)$ is the set of the variables defined by instruction S , and $\text{USE}(S)$ is the set of variables used by instruction S .

Moreover, if the statements are in a loop, a dependence can be loop- independent or loop-
 15 carried. This notion introduces the definition of the distance of a dependence. When a
 dependence is loop- independent ~~it means that~~, it occurs between two instances of different
 statements in the same iteration, and ~~then~~ its distance is equal to 0. ~~On the contrary~~ By
contrast, when a dependence is loop carried, it occurs between two instances in two different
 iterations ~~the dependence is loop carried~~, and ~~the~~ its distance is equal to the difference
 20 between the iteration numbers of the two instances.

The notion of direction of dependence generalizes the notion of distance, and is generally
 used when the distance of a dependence is not constant, or cannot be computed with
 precision. The direction of a dependence is given by $<$, if the dependence between S and S'
 25 occurs when the instance of S is in an iteration before the iteration of the instance of S' , $=$ if
 the two instances are in the same iteration, and $>$ if the instance of S is in an iteration after the
 iteration of the instance of S' .

In the case of a loop nest, ~~we have then~~ there are distance and direction vector, with one
 30 element for each level of the loop nest. ~~The figures below illustrate all these definitions.~~
Figs. 16 to 20 illustrate these definitions. Fig. 16 illustrates a code and diagram of an
example of a true dependence with distance 0 on array 'a'. Fig. 17 illustrates a code and

diagram of an example of an anti-dependence with distance 0 on array 'b'. Fig. 18 illustrates a code and diagram of an example of an output dependence with distance 0 on array 'a'. Fig. 19 illustrates a code and diagram of an example of a dependence with direction vector(=,=) between $S1$ and $S2$ and a dependence with direction vector (= \leq) between $S2$ and $S2$. Fig. 20 illustrates a code and diagram of an example of an anti-dependence with distance vector (0,2).

The data dependence graph ~~is~~may be used by a lot of optimizations, and ~~is~~may also be useful to determine if their application is valid. For instance, a loop can be vectorized if its data dependence graph does not contain any cycle.

```
for (i=0; i<N; i=i+1) {
    S: a[i] = b[i] + 1;
    Si: c[i] = a[i] + 2;
}
```

~~Figure 8: Example of a true dependence with distance 0 on array a~~

```
for (i=0; i<N; i=i+1) {
    S: a[i] = b[i] + 2;
    S1 b[i] = C1[i] + 2;
}
```

~~Figure 9: Example of an anti-dependence with distance 0 on array b~~

```
for (i=0; i<N; i=i+1) {
    S: a[i] = b[i] + 1;
    S1: a[i] = c[i] + 2;
}
```

~~Figure 10: Example of an output dependence with distance 0 on array a~~

```

for (j=0; j<=N; j++)
    for (i=0; i<=N; i++)
    {
5  ----- S1: ----- c[i][j] = 0;

    for (k=0; k<=N; k++)

        ----- S2: ----- c[i][j] = c[i][j] + a[i][k] * b[k][j];

    }

```

~~Figure 11: Example of a dependence with direction vector $(-, -)$ between S1 and S2 and a dependence with direction vector $(-, -, <)$ between S2 and S2.~~

```

for (i=0; i<=N; i++)
    for (j=0; j<=N; j++)

        ----- S: ----- a[i][j] = a[i][j+2] + b[i];

```

~~Figure 12: Example of an anti-dependence with distance vector $(0, 2)$.~~

3.1.3 Interprocedural Alias Analysis

The aim of alias analysis is: to determine if a memory location is aliased by several objects, like e.g., variables or arrays, in a program. It has may have a strong impact on data dependence analysis and on the application of code optimizations. Aliases can occur with statically allocated data, like unions in C where all fields refer to the same memory area, or with dynamically allocated data, which are the usual targets of the analysis. ~~In Figure 13, we have a~~ typical case of aliasing where *p* alias *b*: is:

```

int b [100] , *p;
for (p=b;p < &b[100]; p++)
    *p=0;

```

Figure 13: Example for typical aliasing

5

Alias analysis can be more or less precise depending on whether or not it takes the control-flow into account. When it does, it is called flow-sensitive, and when it does not, it is called flow-insensitive. Flow-sensitive alias analysis is able to detect in which blocks along a path two objects are aliased. As it is more precise, it is more complicated and more expensive to compute. Usually flow-insensitive alias information is sufficient. This aspect is illustrated in Figure 14 in Fig. 21 where a flow-insensitive analysis would find that *p alias b*, but where a flow-sensitive analysis would be able to find that *p alias b* only in block B2.

10

Furthermore, aliases are classified into must-aliases and may-aliases. For instance, if we consider considering flow-insensitive may-alias information, then *x alias y*, iff *x* and *y* may, possibly at different times, refer to the same memory location. And if we consider Considering flow-insensitive must-alias information, *x alias y*, iff *x* and *y* must, throughout the execution of a procedure, refer to the same storage location. In the case of Figure 14, Fig. 21, if we consider flow-insensitive may-alias information is considered, *p alias b* holds, whereas if we consider flow-insensitive must-alias information is considered, *p alias b* does not hold. The kind of information to use depends on the problem to solve. For instance, if we want to remove removal of redundant expressions or statements is desired, must-aliases must be used, whereas if we want to build of a data dependence graph is desired, may-aliases are necessary.

15

20

25

Figure 14: Example of control-flow sensitivity

Finally this analysis must be interprocedural to be able to detect aliases caused by non-local variables and parameter passing. The latter case is depicted in Figure 15 the code below, which is an example for aliasing parameter passing, where *i* and *j* are aliased through the function call where *k* is passed twice as parameter.

30

```

void foo (int *i, int* j)
{
    *i = *j+1;
}
. . .
foo (&k, &k);

```

Figure 15: Example for aliasing by parameter passing

3.1.4—Interprocedural Value Range Analysis

10 This analysis can find the range of values taken by the variables. It can help to apply optimizations like dead code elimination, loop unrolling and others. For this purpose, it can use information on the types of variables and then consider operations applied on these variables during the execution of the program. Thus, it can determine, for instance, if tests in conditional instruction are likely to be ‘met or not, or- determine the iteration range of loop

15 nests.

This analysis has to be interprocedural as, for instance, loop bounds can be passed as parameters of a function, like as in the following example. ~~We know~~ It is known by analyzing the code that in the loop executed with array ‘a’, *N* is at least equal to 11, and that in the loop

20 ~~executed~~ executed with array ‘b’, *N* is at most equal to 10.

```

void foo (int *c, int N)
{
    int i;
    for (i=0; i<N; i++)
        c[i] = g (i, 2);
}
. . .
if (N > 10)
    foo (a, N) ;
else
    foo (b,N) ;

```

The value range analysis can be supported by the programmer by giving further value constraints which cannot be retrieved from the language semantics. This can be done by pragmas or a compiler known assert function.

5 **3.1.5—Alignment Analysis**

Alignment analysis deals with data layout for distributed memory architectures. As stated by Saman Amarasinghe, “Although data memory is logically a linear array of cells, its realization in hardware can be viewed as a multi-dimensional array. Given a dimension in this array, alignment analysis will identify memory locations that always resolve to a single value in that dimension. For example, if the dimension of interest is memory banks, alignment analysis will identify if a memory reference always accesses the same bank.” This is the case in the second part of the figure below, that can be found in [10], where all accesses, depicted in blue Fig. 22, which is a reproduction of a figure that can be found in Sam Larsen, Emmet Witchel & Saman Amarasinghe, “Increasing and Detecting Memory Address Congruence,” *Proceedings of the 2002 IEEE International Conference on Parallel Architectures and Compilation Techniques (PACT'02)*, 18-29 (September 2002). All accesses, depicted in dark squares, occur to the same memory bank, whereas in the first part, the accesses are not aligned. He Saman Amarasinghe adds then that: “Alignment information is useful in a variety of compiler-controlled memory optimizations leading to improvements in programmability, performance, and energy consumption.”

Alignment analysis, for instance, is able to help find a good distribution scheme of the data and is furthermore useful for automatic data distribution tools. An automatic alignment analysis tool can be able to automatically generate alignment proposals for the arrays accessed in a procedure and thus simplifies the data distribution problem. This can be extended with an interprocedural analysis taking into account dynamic realignment.

Alignment analysis can also be used to apply loop alignment that transforms the code directly rather than the data layout in itself, as shown later. discussed below. Another solution can be used for the PACT XPP, relying on the fact that it can handle aligned code very efficiently. It ~~consists in~~ includes adding a conditional instruction testing if the accesses in the loop body are aligned followed by the necessary number of peeled iterations of the loop body, then the aligned loop body, and then some compensation code. Only the aligned code is then executed by the PACT XPP, ~~the~~ The rest is may be executed by the host processor. If the

alignment analysis is more precise (inter-procedural or inter-modular), less conditional code has to be inserted.

3.2—Code Optimizations

5 ~~Most~~Discussion regarding many of the optimizations and transformations ~~presented~~
~~here~~discussed below can be found in detail in [4], ~~and also in [2,3,5].~~David F. Bacon, Susan
L. Graham & Oliver J. Sharp, “Compiler Transformations for High-Performance
Computing,” ACM Computing Surveys, 26(4):325-420 (1994); Michael Wolfe, *supra*; Hans
Zima et al., *supra*; and Steven Muchnick, *supra*.

3.2.1—General Transformations

~~We present in this section~~Discussed below are a few general optimizations that can be
applied to straightforward code, and to loop bodies. These are not the only ones that appear
in a compiler, ~~but they are mentioned in the sequel of this document.~~

Constant Propagation

~~This optimization propagates~~A constant propagation may propagate the values of constants
into the expressions using them throughout the program. This way a lot of computations can
be done statically by the compiler, leaving less work to be done during the execution, ~~this.~~

20 This part of the optimization ~~is~~is also known as constant folding.

An example of constant propagation is:

```

N = 256;                                for(i=0; i<=256; i++)
c = 3;                                  a[i] = b[i] + 3;
25  for (i=0; i<=N; i++)
      a[i] = b[i] + c;
```

Figure 16: Example of constant propagation

Copy Propagation

30 ~~This~~A copy propagation optimization simplifiesmay simplify the code by removing
redundant copies of the same variable in the code. These copies- can be produced by the
programmer ~~himself~~- or by other optimizations. This optimization ~~reduces~~may reduce the
register pressure and the number of register-to-register move instructions.

An example of copy propagation is:

```

t = i*4;
r = t;
5  for (i=0; i<=N; i++)
    a[r] = b[r] + a[i];

t = i*4;-
for (i=0; i<=N; i++)
    a[t] = b[t] + a[i];
```

Figure 17: Example of copy propagation

Dead Code Elimination

10 ~~This~~A dead code elimination optimization ~~removes~~may remove pieces of code that will never be executed. Code is never executed if it is in the branch of a conditional statement whose condition is always evaluated to true or false, or if it is a loop body, whose number of iterations is always equal to 0.

15 Code updating variables, that are never used, is also useless and can be removed as well. If a variable is never used, then the code updating it and its declaration can also be eliminated.

An example of dead code elimination is:

```

20  for (i=0; i<=N; i++) {
    for (j=0; j<0; j++)
        a[j] = b[j] + a[i];
    for (j=0; j<10; j++)
        a[j+1] = a[j] + b[j];
    }

for (i=0; i<=N; i++) {
    for (j=0; j<10; j++)
        a[j+1] = a[j] + b[j];
}
```

25 ~~Figure 18: Example of dead code elimination~~

Forward Substitution

30 ~~This~~A forward substitution optimization is a generalization of copy propagation. The use of a variable ~~is~~may be replaced by its defining expression. It can be used for simplifying the data dependency analysis and- the application of other transformations by making the use of loop variables visible.

An example of forward substitution is:

```
c = N + 1;                                for (i=0; i<=N; i++)
for (i=0; i<= N; i++)                      a[N+1] = b[N+1] + a[i];
a[c] = b[c] + a[i];
```

Figure 19: Example of forward substitution

Idiom Recognition

ThisAn idiom recognition transformation recognizesmay recognize pieces of code and can replace them by calls to compiler known functions, or less expensive code sequences, like code for absolute value computation.

An example of idiom recognition is:

```
for (i=0; i<N; i++) {                      for (i=0; i<N; i++) {
    c = a[i] - b[i];                        c = a[i] - b[i];
    if (c<0)                                c = abs (c) ;
        c = -c;                             d[i] = c;
    d[i] = c;                               }
}
```

Figure 20: Example of idiom recognition

3.2.2—Loop Transformations

Loop Normalization

ThisA loop normalization transformation ensuresmay ensure that the iteration space of the loop is always with a lower bound equal to 0 or 1 (depending on the input language), and with a step of 1. The array subscript expressions and the bounds of the loops are modified accordingly. It can be used before loop fusion to find opportunities, and ease inter-loop dependence analysis, and it also enables the use of dependence tests that ~~needs~~need a normalized loop to be applied:

An example of loop normalization is:

```
for (i=2; i<N; i=i+2)                      for (i=0; i<(N-2)/2; i++)
    a[i] = b[i];                            a[2*i+2] = b[2*i+2];
```

Figure 21: Example of loop normalization

5 **Loop Reversal**

~~This~~A loop reversal transformation changes~~may change~~ the direction in which the iteration space of a loop is scanned. It is usually used in conjunction with loop normalization and other transformations, like loop interchange, because it changes the dependence vectors.

10 An example of loop reversal is:

for (i=N; i>=0; i--)	for (i=0; i<=N; i++)
a[i] = b[i];	a[i] = b[i];

Figure 22: Example of loop reversal

15 **Strength Reduction**

~~This~~A strength reduction transformation replaces~~may replace~~ expressions in the loop body by equivalent but less expensive ones. It can be used on induction variables, other than the loop variable, to be able to eliminate them.

20 An example of strength reduction is :

for (i=0; i<N; i++)	t = c;
a[i] = b[i] + c*i;	for (i=0; i<N; i++) {
	a[i] = b[i] + t;
	t = t + c;
	}

25

Figure 23: Example of strength reduction

Induction Variable Elimination

~~This~~An induction variable elimination transformation can use strength reduction to remove
induction variables from a loop, hence reducing the number of computations and easing the
analysis of the loop. This may also remove ~~removes~~ dependence cycles due to the update of
the variable, enabling vectorization.

An example of induction variable elimination is:

```

5      for (i=0 ; i<=N; i++)      {      for (i=0; i<=N; i++){
      k = k + 3;                    a[i] = b[i] + a[k+(i+1)*3];
      a[i] = b[i] + a[k];      }
      }

      k = k + (N+1) *3;

```

Figure 24: Example of induction variable elimination

Loop-invariant Invariant Code Motion

10 This A loop-invariant code motion transformation ~~moves~~may move computations outside a loop if their result is the same in all iterations. This ~~allows to reduce~~may allow a reduction of the number of computations in the loop body. This optimization can also be conducted in the reverse fashion in order to get perfectly nested loops, that are easier to handle by other optimizations.

15

An example of loop-invariant code motion is:

```

      for (i=0 ; i<N; i++)      if (N >= 0)
      a [i] = b[i] + x*y;        c = x*y;
                                for (i=0; i<N; i++)
                                a [i] = b [i] + c;
20

```

Figure 25: Example of loop-invariant code motion

Loop Unswitching

25 This A loop unswitching transformation ~~moves~~may move a conditional instruction outside of a loop body if its condition is loop- invariant. The branches of the condition ~~are~~may then be made of the original loop with the appropriate original statements of the conditional statement. It ~~allows~~may allow further parallelization of the loop by removing control-flow in the loop body and also removing unnecessary computations from it.

30

An example of loop unswitching is:

```

    for (i =0; i<N; i++) {
        a[i] = b[i] + 3;
        if (x > 2)
            b[i] = c[i] + 2;
        else
            b[i]=c[i] - 2;
    }

    if (x > 2)
        for (i=0; i<N; i++){
            a[i] = b[i] + 3;
            b[i] = c[i] +2;
        }
    else
        for (i=0; i<N; i++){
            a[i] = b[i] + 3;
            b[i] = c[i]- 2;
        }

```

Figure 26: Example of loop unswitching

If-Conversion

15 This An if-conversion transformation is may be applied on loop bodies with conditional instructions. It changes may change control dependencies dependencies into data dependencies dependencies and allows allow then vectorization to take place. It can be used in conjunction conjunction with loop unswitching to handle loop bodies with several basic blocks. The conditions, where array expressions could appear, are may be replaced by Boolean boolean terms called guards. Processors with predicated execution support can execute directly such code.

An example of if-conversion is:

```

    for (i=0; i<N; i++){
        a[i] = a[i] + b[i];
        if (a[i] != 0)
            if (a[i] > c[i])
                a[i] = a[i] - 2;
            else
                a[i] = a[i] + 1;
        d[i] = a[i] * 2;
    }

    for (i=0; i<N; i++){
        a[i] = a[i] + b[i];
        c2 = (a[i] != 0);
        if (c2) c4 = (a[i] > c[i]);
        if (c2 && c4) a[i] = a[i] - 2;
        if (c2 && ! c4) a[i] = a[i] + 1;
        d[i] = a[i] * 2;
    }

```

Figure 27: Example of if-conversion

Strip-Mining

~~This~~A strip-mining transformation ~~enables to adjust~~may enable adjustment of the granularity of an operation. It is commonly used to choose the number of independent computations in the inner loop nest. When the iteration count is not known at compile time, it can be used to generate a fixed iteration count inner loop satisfying the resource constraints. It can be used in conjunction with other transformations like loop distribution or loop interchange. It is also called loop sectioning. Cycle shrinking, also called stripping, is a specialization of strip-mining.

An example of strip-mining is:

```
for (i=0; i<N; i++)      up = (N/16) *16;
    a[i] = b[i] + C;      for(i=0; i<up; i = i + 16)
                           a[i:1+16] = b[i:i+16] + c;
                           for (j =i+1; j <N; j++)
                               a[i] = b[i] + c;
```

Figure 28: Example of strip-mining

Loop Tiling

~~This~~A loop tiling transformation ~~modifies~~may modify the iteration space of a loop nest by introducing loop levels to divide the iteration space in tiles. It is a multi-dimensional generalization of strip-mining. It is generally used to improve memory reuse, but can also improve processor, register, TLB, or page locality. It is also called loop blocking.

The size of the tiles of the iteration space ~~is~~may be chosen so that the data needed in each tile fit in the cache memory, thus reducing the cache misses. In the case of coarse-grain computers, the size of the tiles can also be chosen so that the number of parallel operations of the loop body ~~fit~~fits the number of processors of the computer.

An example of loop tiling is:

```

for (i=0; i<N; i++)      for (ii=0; ii<N; ii = ii+16)
    for (j=0; j<N; j++)    for (jj=0; jj<N; jj = jj+16)
        a[i][j] = b[j][i];    for (i=ii; i<min(ii+15,N); j++)
5          for (j=jj; j<min(jj+15,N); j++)
                        a[i][j] = b[j][i];

```

Figure 29: Example of loop tiling

Loop Interchange

- 10 This A loop interchange transformation ~~is~~may be applied to a loop nest to move inside or outside (depending on the searched effect) the loop level containing data ~~dependences-dependencies~~. It can:
- ~~■~~—enable vectorization by moving inside an independent loop and outside a dependent loop, or
 - 15 • ~~■~~—improve vectorization by moving inside the independent loop with the largest range, or
 - ~~■~~—deduce the stride, or
 - ~~■~~—increase the number of loop-invariant expressions in the inner-loop, or
 - ~~■~~—improve parallel performance by moving an independent loop outside of a loop nest
 - 20 to increase the granularity of each iteration and reduce the number of barrier synchronizations.

An example of a loop interchange is:

```

for (i=0; i<N; i++)      for (j=0; j<N; j++)
25   for (j=0; j<N; j++)    for (i=0; i<N; i++)
        a[i] = a[i] + b[i][j];    a[i] = a[i] + b[i][j];

```

Figure 30: Example of loop interchange

Loop Coalescing / Collapsing

- 30 This A loop coalescing / collapsing transformation ~~combines~~may combine a loop nest into a single loop. It can improve the scheduling of the loop, and also reduces the loop overhead. Collapsing is a simpler version of coalescing in which the number of dimensions of arrays is reduced as well. Collapsing ~~reduces~~may reduce the overhead of nested loops and ~~multi-~~

~~dimensional~~multidimensional arrays. Collapsing can be applied to loop nests that iterate over memory with ~~at a~~ constant stride. ~~Otherwise, otherwise~~-loop coalescing ~~is~~may be a better approach. It can be used to make vectorizing profitable by increasing the iteration range of the innermost loop.

5

An example of loop coalescing is:

```

for (i=0; i<N; i++)          for (k=0; k<N*M; k++) {
    for (j=0; j<M; j++)      i = ((k-1) /m) *m + 1;
        a[i][j] = a[i][j] + c; j = ((T-1) %m) + 1;
10                        a[i][j] = a[i][j] + c;
                                }

```

Figure 31: Example of loop coalescing

Loop Fusion

15 ~~This~~A loop fusion transformation, also called loop jamming, ~~merges 2~~ may merge two successive loops. It ~~reduces~~may reduce loop overhead, increases instruction-level parallelism, improves register, cache, TLB or page locality, and improves the load balance of parallel loops. Alignment can be taken into account by introducing conditional instructions to take care of ~~dependences~~dependencies.

20

An example of loop fusion is:

```

for (i=0; i<N; i++)          for (i=0; i<N; i++) {
    a[i] = b[i] + c;          a[i] = b[i] + c;
                                d[i] = e[i] + c;
25 for (i=0; i<N; i++)      }
                                d[i] = e[i] + c;

```

Figure 32: Example of loop fusion

Loop Distribution

30 ~~This~~A loop distribution transformation, also called loop fission, ~~allows~~may allow to split a loop in several pieces in case the loop body is too big, or because of

~~dependenciesdependencies.~~ The iteration space of the new loops ~~is~~may be the same as the iteration space of the original loop. Loop spreading is a more sophisticated distribution.

An example of loop distribution is:

```

5      for (i=0; i<N; i++)          {   for (i=0; i<N; i++)
      a[i] = b[i] + c;                a[i] = b[i] + c;
      d[i] = e[i] + c;
    }                                for (i=0; i<N; i++)
                                    d[i] = e[i] + c;

```

Figure 33: Example of loop distribution

Loop Unrolling / Unroll-and-Jam

~~This~~A loop unrolling / unroll-and-jam transformation ~~replicates~~may replicate the original loop body in order to get a larger one. A loop can be unrolled partially or completely. It ~~is~~may be used to get more opportunity for parallelization by making the loop body bigger, ~~it~~. It ~~may~~ also ~~improves~~improve register, or cache usage and reduces loop overhead. Loop unrolling the outer loop followed by merging the induced inner loops is referred to as unroll-and-jam.

An example of loop unrolling is:

```

20      for (i=0; i<N; i++)          for (i=0; i<N; i = i+2) {
      a [i] = b [i] + c;                a[i] = b[i] + c;
                                      a[i+1] = b[i+1] + c;
                                      }
25      if ((N-1)%2) == 1
                                      a[N-1 ] = b [N-1] + c;

```

Figure 34: Example of loop unrolling

Loop Alignment

~~This~~A loop alignment optimization ~~transforms~~may transform the code to get aligned array accesses in the loop body. Its effect ~~it~~may be to transform loop-carried ~~dependenciesdependencies~~ into loop-independent ~~dependenciesdependencies~~, which allows to ~~extract~~for extraction of more parallelism from a loop. It can use different transformations,

like loop peeling or introduce conditional statements, to achieve its goal. This transformation can be used in conjunction with loop fusion to enable this optimization by aligning the array accesses in both loop nests. In the example below, all accesses to array 'a' become aligned.

5 An example of loop alignment is:

```

for (i=2; i <= N; i++){          for (i=1; i<=N; i++){
    a[i] = b[i] + c[i];           if (i>1) a[i] = b[i] + c[i];
    d[i] = a[i-1] * 2;           if (i<N) d[i+1] = a[i] * 2;
    e[i] = a[i-1] + d[i+1];       if (i<N) e[i+1] = a[i] + d[i+2];
10 }                             }

```

Figure 35: Example of loop alignment

Loop Skewing

15 ThisA loop skewing transformation ismay be used to enable parallelization of a loop nest. It ismay be useful in combination with loop interchange. It ismay be performed by adding the outer loop index multiplied by a skew factor, f , to the bounds of the inner loop variable, and then subtracting the same quantity from every use of the inner loop variable inside the loop.

An example of loop skewing is:

```

20 for (i=1; i<=N; i++){          for (i=1; i<=N; i++){
    for (j=1; j<=N; j++)           for (j=i+1; j<=i+N; j++)
        a[i] = a[i+j] + c;         a[i] = a[j] + c;

```

Figure 36: Example of loop skewing

25 Loop Peeling

ThisA loop peeling transformation removesmay remove a small number of beginning or ending iterations of a loop to avoid dependences in the loop body. These removed iterations aremay be executed separately. It can be used for matching the iteration control of adjacent loops to enable loop fusion.

30

An example of loop peeling is:

```

for (i=0; i<=N; i++)          a[0][N] = a[0][N] + a[N][N];
    a[i][N] = a[0][N] + a[N][N];  for (i=1; i<=N-1; i++)
                                a[i][N] = a[0][N] + a[N][N];
5                                a[N][N] = a[0][N] + a[N][N];

```

~~Figure 37: Example of loop peeling~~

Loop Splitting

~~This~~ A loop splitting transformation ~~cuts~~may cut the iteration space in pieces by creating other loop nests. It is also called Index Set Splitting, and is generally used because of ~~dependencies~~dependencies that prevent parallelization. The iteration space of the new loops ~~is~~may be a subset of the original one. It can be seen as a generalization of loop peeling.

An example of loop splitting is:

```

15 for (i=0; i<=N; i++)          for (i=0; i<(N+1)/2; i++)
    a[i] = a[N-i+1] + c;          a[i] = a[N-i+1] + c;
                                for (i = (N+1)/2; i<=N; i++)
                                a[i] = a[N-i+1] + c;

```

~~Figure 38: Example of loop splitting~~

Node Splitting

~~This~~ A node splitting transformation ~~splits~~may split a statement in pieces. It ~~is~~may be used to break dependence cycles in the dependence graph due to the too high granularity of the nodes, thus enabling vectorization of the statements.

An example of node splitting is:

```

for (i=0; i<N; i++) {          for (i=0; i<N; i++) {
    b[i] = a[i] + c[i] * d[i];    t1[i] = c[i] * d[i];
    a[i+1] = b[i] * (d[i] - c[i]); t2[i] = d[i] - c[i];
30 }                             b[i] = a[i] + t1[i];
                                a[i+1] = b[i] * t2[i];
                                }

```

Figure 39: Example of node splitting

Scalar Expansion

This scalar expansion transformation ~~replaces~~may replace a scalar in a loop by an array to eliminate ~~dependences~~dependencies in the loop body and enable parallelization of the loop nest. If the scalar is used after the loop, a compensation code must be added.

An example of scalar expansion is:

```
10  for (i=0; i<N; i++) {           for (i=0; i<N; i++) {
    c = b[i];   tmp[i] = b[i];      a[i] = a[i] + tmp[i];
    a[i] = a[i] + c;               }
    }                               c = tmp[N-1];
```

Figure 40: Example of scalar expansion

Array Contraction / Array Shrinking

This array contraction / array shrinking transformation is the reverse transformation of scalar expansion. It may be needed if scalar expansion generates too many memory requirements.

An example of array contraction is:

```
25  for (i=0; i<N; i++)           for (i=0; i<N; i++)
    for (j=0; j<N; j++){         for (j=0; j<N; j++){
        t[i][j] = a[i][j] * 3;    t[j] = a[i][j] * 3;
        b[i][j] = t[i][j] + c[j]; b[i][j] = t[j] + c[j];
    }                             }
```

Figure 41: Example of array contraction

Scalar Replacement

30 This scalar replacement transformation ~~replaces~~may replace an invariant array reference in a loop, by a scalar. This array element ~~is~~may be loaded in a scalar before the inner loop and

stored again after the inner loop, if it is modified. It can be used in conjunction with loop interchange.

An example of scalar replacement is:

<pre> for (i=0; i<N; i++) 5 for (j=0; j<N; j++) a[i] = a[i] + b[i][j]; </pre>	<pre> for (i=0; i<N; i++) { tmp = a[i]; for (j=0; j<N; j++) tmp = tmp + b[i][j]; a[i] = tmp; } </pre>
---	---

10

Figure 42: Example of scalar replacement

Reduction Recognition

~~This~~ A reduction recognition transformation allows to handle ~~may allow handling of~~ reductions in loops. A reduction ~~is~~ may be an operation that computes a scalar value from arrays. It can be a dot product, the sum or minimum of a vector for instance. ~~The~~ A goal is then to perform as many operations in parallel as possible. One way ~~is~~ may be to accumulate a vector register of partial results and then reduce it to a scalar with a sequential loop. Maximum parallelism ~~is~~ may then be achieved by reducing the vector register with a tree; i.e., pairs of elements ~~elements~~ are summed; then pairs of these results are summed; etc.

20

An example of reduction recognition is:

<pre> for (i=0; i<N; i++) s = s + a[i]; </pre>	<pre> for (i=0; i<N; i=i+64) tmp[0:63] = tmp[0:63] + a[i:i+63]; for (i=0; i<64; i++) 25 s = s + tmp[i]; </pre>
---	--

25

Figure 43: Example of reduction recognition

Loop Pushing / Loop Embedding

~~This~~ A loop pushing / loop embedding transformation replaces ~~may replace~~ a call in a loop body by the loop in the called function. It ~~is~~ may be an ~~inter-procedural~~ interprocedural optimization. It ~~allows~~ may allow the parallelization of the loop nest and ~~eliminates~~ eliminate

30

the overhead caused by the procedure call. Loop distribution can be used in conjunction with loop pushing.

An example of loop pushing is:

```

5  for (i=0; i<N; i++)      f2 (x)
    f(x,i);

                                void f2 (int* a) {
void f (int* a, int j) {      for (i=0; i<N; i++)
    a[j] = a[j] + c;          a[i] = a[i] + c;
10 }                          }

```

Figure 44: Example of loop pushing

Procedure Inlining

This procedure inlining transformation replaces a call to a procedure by the code of the procedure itself. It is an ~~inter-procedural~~ interprocedural optimization. It allows a loop nest to be parallelized, removes overhead caused by the procedure call, and can improve locality.

An example of procedure inlining is:

```

20  for (i=0; i<N; i++)      for (i=0; i<N; i++)
    f(a,i);                  a[i] = a[i] + c;

void f (int* x, int j) {
    x[j] = x[j] + c;
}

```

Figure 45: Example of procedure inlining

Statement Reordering

This statement reordering transformation schedules instructions of the loop body to modify the data dependence graph and enable vectorization.

An example of statement reordering is:

```

5      for (i=0; i<N; i++) {      for(i=0; i<N; i++){
          a[i] = b[i] * 2;          c[i] = a[i-1] - 4;
          c[i] = a[i-1] - 4;      a[i] = b[i] * 2;
                                  }
                                  }

```

Figure 46: Example of statement reordering

Software Pipelining

10 This A software pipelining transformation parallelizes ~~may parallelize~~ a loop body by scheduling instructions of different instances of the loop body. It ~~is~~ may be a powerful optimization to improve instruction-level parallelism. It can be used in conjunction with loop unrolling. In the example below, -the preload commands can be issued one after another, each taking only one cycle. This time is just enough to request the memory areas. It is not enough to actually load them. This takes many cycles, depending on the cache level that

15 actually has the data. Execution of a configuration behaves similarly. The configuration is issued in a single cycle, waiting until all data are present. Then the configuration executes for many cycles. Software pipelining overlaps the execution of a configuration with the preloads for the next configuration. This way, the XPP array can be kept busy in parallel to the Load/Store unit.

20

An example of software pipelining is:

```

Issue Cycle Command
      XPPPreloadConfig(CFG1) ;
      for (i=0; i<100; ++i) {
25      1:      XPPPreload (2, a+10* i,10) ;
          2:      XPPPreload (5,b+20*i, 20) ;
          3:
          4:      // delay
          5:
30      6:      XPPEecute (CFG1);
          }

```

Issue Cycle Command


```

Prologue XPPPreloadConfig(CFG1);
        XPPPreload (2, a,10) ;
        XPPPreload (5, b,20) ;
        // delay
5         for (i=1; i<100; ++i) {
Kernel
        1:      XPPEXecute (CFG 1) ;
        2:      XPPPreload (2, a +10*i,10) ;
        3:      XPPPreload (5, b+20*i, 20);
10       4:      }
        XPPEXecute (CFG1) ;
Epilog // delay

```

Figure 47: Example of software pipelining

15 **Vector Statement Generation**

This A vector statement generation transformation ~~replaces~~may replace instructions by vector instructions that can perform an operation on several data in parallel.

An example of vector statement generation is:

```

20         for (i=0; i<N; i++)           [0:N] = b[0:N];
        [i] = b[i];

```

Figure 48: Example of vector statement generation

25 **3.2.3—Data-Layout Optimizations**

~~In the following we describe optimizations that~~Optimizations may modify the data layout in memory in order to extract more parallelism or prevent memory problems like cache misses. Examples of such optimizations are scalar privatization, array privatization, and array merging.

30 **Scalar Privatization**

~~This~~A scalar privatization optimization ~~is~~may be used in multi-processor systems to increase the amount of parallelism and avoid unnecessary communications between the processing elements. If a scalar is only used like a temporary variable in a loop body, then each

processing element can receive a copy of it and achieve its computations with this private copy.

An example of scalar privatization is:

```

5      for (i=0; i <= N; i++) {
          c = b[i];
          a[i] = a[i] + c;
      }

```

Figure 49: Example for scalar privatization

10

Array Privatization

~~This~~An array privatization optimization ~~is~~may be the same as scalar privatization except that it ~~works~~may work on arrays rather than on scalars.

15

Array Merging

~~This~~An array merging optimization ~~transforms~~may transform the data layout of arrays by merging the data of several arrays following the way they are accessed in a loop nest. This way, memory cache misses can be avoided. The layout of the arrays can be different for each loop nest. ~~Below is the~~The example code for array merging presented below is an example of a cross-filter, where the accesses to ~~may~~array 'a' are interleaved with accesses to array ~~b~~. ~~The picture next to it represents the data lay-out~~'b'. Fig. 23 illustrates a data layout of both arrays, where blocks of ~~a~~ ~~(in green)~~'a' (the dark highlighted portions) are merged with blocks of ~~b~~ ~~(in yellow)~~'b' (the lighter highlighted portions). Unused memory space is ~~in~~represented by the white portions. Thus, cache misses ~~are~~may be avoided as data blocks containing arrays 'a' and 'b' are loaded into the cache when getting data from memory. More details can be found in [11]Daniela Genius & Sylvain Lelait, "A Case for Array Merging in Memory Hierarchies," Proceedings of the 9th International Workshop on Compilers for Parallel Computers, CPC'01 (June 2001).

```

25      for (j=1; j<=N-1; i++)
30      for (j=1; j<=N; j++)
          b[i][j] = 0.25*(a[i-1][j]+a[i][j-1]+a[i+1][j]+a[i][j+1]);

```

Figure 50: Example for array merging

3.2.4—Example of application of the optimizations

~~As seen before~~In accordance with that which is discussed above, it will be appreciated that a
lot of optimizations can be performed on loops before and also after generation of vector

statements. Finding a sequence of optimizations that- would produce an optimal solution for

5 all loop nests of a program is still an area of research. Therefore ~~we can only propose,~~ in an
embodiment of the present invention, a way to use these optimizations is provided that

follows a reasonable heuristic to produce vectorizable loop nests. To vectorize the code, ~~we~~
~~can use the Allen-Kennedy algorithm,~~ that uses statement reordering and loop distribution
before vector statements are generated-, can be used. It can be enhanced with loop

10 interchange, scalar expansion, -index set splitting, node splitting, loop peeling. All these
transformations are based on the data dependence graph. A statement can be vectorized if it
is not part of a dependence cycle. Hence, ~~hence~~ optimizations are may be performed to break
cycles or, if not completely possible, to create loop nests without dependence cycles.

15 ~~We can divide the~~ The whole process in may be divided into four majors steps. First ~~we~~
~~should restructure,~~ the procedures may be restructured by analyzing the procedure calls inside
the loop bodies ~~and try to remove them. Then,~~ Removal of the procedures may then be tried.
Then, some high-level dataflow optimizations are may be applied to the loop bodies to modify
their control-flow and simplify their code. The third step ~~would consist in~~ may include

20 preparing the loop nests for vectorization by building perfect loop nests and ensuring that
inner loop levels are vectorizable. Then, optimizations can be performed that target the
architecture and optimize the data locality. It should also be noted that other optimizations
and code transformations can occur between these different steps that can also help to further
optimize the loop nests.

25 Hence, the first step ~~applies~~ may apply procedure inlining and loop pushing to remove the
procedure calls of the loop bodies. Then, the second step ~~consists of~~ may include loop-
invariant code motion, loop unswitching, strength reduction and idiom recognition. The third
step can be divided in several subsets of optimizations. ~~We can first apply loop~~ Loop
30 reversal,- loop normalization and if-conversion may be initially applied to get normalized
loop nests. This ~~allows to build~~ may allow building of the data dependency graph. Then, if
~~dependences~~ dependencies prevent the loop nest to be vectorized, transformations ~~can~~ may be
applied. For instance, if ~~dependences~~ dependencies occur only on certain iterations, loop
peeling or loop splitting ~~can~~ may be applied. Node splitting, loop skewing, scalar expansion

or statement reordering can be applied in other cases. Then, loop interchange ~~moves~~may
move inwards the loop levels without dependence cycles. ~~The~~A goal is to have perfectly
nested loops with the loop levels carrying dependence cycles as much outwards as possible.

~~Then we can apply,~~ loop, fusion, reduction recognition, scalar replacement / array

- 5 contraction, and loop distribution may be applied to further improve the following
vectorization. Vector statement generation can be performed at last using the Allen-Kennedy
algorithm for instance. The last step can ~~consist of~~include optimizations ~~like~~such as loop
tiling, strip-mining, loop unrolling and software pipelining that take into account the target
processor.

10

The number of optimizations in the third step ~~is~~may be large, but it may be that not all of
them are applied to each loop nest. Following the goal of the vectorization and the data
dependence graph, only some of them are applied. Heuristics ~~are~~may be used to guide the
application of the optimizations, that can be applied several times if needed. ~~Let us illustrate~~

- 15 ~~this with an example.~~ The following code is an example of this:

```
void f (int** a, int** b, int *c, int i, int j) {  
    a[i] [j] = a[i] [j-1] - b[i+1][j-1];  
}
```

20

```
void g(int* a , int* c, int i){  
    a[i] = c[i] + 2;  
}
```

25

```
    for (i=0; i<N; i++){  
        for (j=1; j<9; j=j++){  
            if (k>0)  
                f(a, b, i, j);  
            else  
                g(d, c, j);
```

```

    }

    d[i] = d[i+1] + 2;

}

for (i=0; i<N; i++)
5      a[i][i] = b[i] + 3;

```

The first step will find that inlining the two procedure calls is possible, ~~then.~~ Then loop unswitching can be applied to remove the conditional instruction of the loop body. The second step ~~begins~~ may begin by applying loop normalization and analyses of the data
 10 dependence graph. A cycle can be broken by applying loop interchange as it is only carried by the second level. The two levels ~~are~~ may be exchanged, so that the inner level is vectorizable. Before ~~that or also after,~~ we apply loop distribution ~~may be applied.~~ Loop fusion can be applied when the loop on *i* is pulled out of the conditional instruction by a traditional redundant code elimination optimization. Finally, vector code can be generated
 15 for the resulting loops.

~~So in~~ In more details ~~detail,~~ after procedure inlining, ~~we obtain~~ the following may be obtained:

```

    for (i=0; i<N; i++) {
        for (j =1; j <9; j=j++)
            if (k>0)
20              a[i][j] = a[i][j-1] - b[i+1][j-1];
            else
                d[j] = c[j] + 2;
        }
25      d[i] = d[i+1] + 2;
    }

    for (i==0 ; i<N; i++)
        a[i][i] = b [i] + 3;
30

```

After loop unswitching, ~~we obtain~~the following may be obtained:

```
    if (k > 0)
        for (i=0; i<N; i++){
            for (j=1; j<9; j=j++)
                a[i][j] = a[i][j-1] - b[i+1][j-1];
5          d[i] = d[i+1] + 2;
            }
        else
            for (i=0; i<N; i++){
10          for (j=1; j<9; j=j++)
                d[j] = c[j] + 2;
                d[i] = d[i+1] + 2;
            }
.

15    for (i=0 ; i<N; i++)
        a [i] [i] = b[i] + 3;
```

After loop normalization, ~~we obtain~~the following may be obtained:

```
    if (k > 0)
20    for (i=0; i<N; i++) {
        for (j=0; j<8; j=j++)
            a[i] [j+1] = a[i] [j] - b [ i+1] [j];
            d[i] = d[i+1] + 2;
        }
25    else
        for (i=0; i<N; i++){
            for (j=0; j<8; j = j++)
                d[j]  = c[j+1] + 2;
```

```

        d[i] = d[i+1] + 2;

    }

```

```

    for (i=0; i<N; i++)

5        a[i][i] = b[i] + 3;

```

After loop distribution and loop fusion, ~~we obtain~~the following may be obtained:

```

    if (k > 0)
        for (i=0; i<N; i++)
10        for (j=0; j<8; j=j++)
            a[i][j+1] = a[i][j] - b[i+1][j];
    else
        for (i=0; i<N; i++)
            for (j=0; j<8; j=j++)
15        d[j] = c[j+1] + 2;

    for (i=0; i<N; i++) {
        d[i] = d[i+1] + 2;
        a[i][i] = b[i] + 3;
20    }

```

After loop interchange, ~~we obtain~~the following may be obtained:

```

    if (k > 0)
        for (j=0; j<8; j=j++)
25        for (i=0; i<N; i++)
            a[i][j+1] = a[i][j] - b[i+1][j];
    else
        for- (i=0; i<N; i++)
            for (j=0; j<8; j=j++)
30        d[j] = c[j+1] + 2;

    for (i=0; i<N; i++) {

```

```

    d[i] = d[i+1] + 2;
    a[i][i] = b[i] + 3;
}

```

5 After vector code generation, ~~we obtain~~ the following may be obtained:

```

    if (k > 0)
        for (j=0; j < 8; j=j++)
            a[0:N-1][j+1] = a[0:N--1][j] - b[0:N][j];
    else
10        for (i=0; i < N; i++)
            d[0:8] = c[1:9] + 2;

    d[0:N-1] = d[1:N] + 2;
    a[0:N-1][0:N-1] = b[0:N] + 3;

```

15

4 ~~COMPILER SPECIFICATION FOR THE PACT XPP~~

4.1 ~~Introduction~~

A cached RISC-XPP architecture ~~exploits~~ may exploit its full potential on code that is characterized by high data locality and high computational effort. A compiler for this
20 architecture has to consider these design constraints. The compiler's primary objective is to concentrate computational expensive calculations to innermost loops and to make up as much data locality as possible for them.

The compiler ~~contains~~ may contain usual analysis and optimizations. As interprocedural
25 analysis, ~~like e.g.,~~ alias analysis₁, are especially useful, a global optimization driver ~~is~~ may be necessary to ensure the propagation of global information to all optimizations. The way the PACT XPP may influence the compiler is discussed in the following sections ~~concentrate on the way the PACT XPP influences the compiler.~~

30 4.2 ~~Compiler Structure~~

~~Figure 51~~ Fig. 24 provides a global view of the compiling procedure and shows the main steps
the compiler ~~must~~ may follow to produce code for a system containing a RISC processor and

a PACT XPP. The next sections focus on the XPP compiler itself, but first the other steps are briefly described.

~~Figure 51: Global View of the Compiling Process~~

5 **4.2.1—Code Preparation**

This ~~step takes~~ Code preparation may take the whole program as input and can be considered as a usual compiler front-end. It ~~will~~ may prepare the code by applying code analysis and optimizations to enable the compiler to extract as many loop nests as possible to be executed by the PACT XPP. Important optimizations are idiom recognition, copy propagation, dead
10 code elimination, and all usual analysis like dataflow and alias analysis.

4.2.2—Partitioning

Partitioning ~~decides~~ may decide which part of the program is executed by the host processor and which part is executed by the PACT XPP.

15

A loop nest ~~is~~ may be executed by the host in three cases:

- ~~—~~if the loop nest is not well-formed,
- ~~—~~if the number of operations to execute is not worth ~~it to be~~ being executed on the PACT XPP, or
- 20 • ~~—~~if it is impossible to get a mapping of the loop nest on the PACT XPP.

A loop nest is said to be well-formed if the loop bounds and the step of all loops are constant, the loop induction variables are known and if there is only one entry and one exit to the loop nest.

25

Another problem ~~arises~~ may arise with loop nests where the loop bounds are constant but unknown at compile time. Loop tiling ~~allows to overcome~~ may allow for overcoming this problem, ~~it as~~ will be described below. Nevertheless, ~~it could be that it is not worth it to execute~~ executing the loop nest on the PACT XPP if the loop bounds are too low. A
30 conditional instruction testing if the loop bounds are large enough can be introduced, and ~~2 two~~ versions of the loop nest ~~are~~ may be produced. One would be executed on the host processor, and the other on the PACT XPP when the loop bounds are suitable. This would

also ease applications of loop transformations, as possible compensation code would be simpler due to the hypothesis on the loop bounds.

4.2.3—RISC Code Generation and Scheduling

5 After the XPP compiler has produced NML code for the loops chosen by the partitioning phase, the main compiling process ~~must~~may handle the code that will be executed by the host processor where instructions to manage the configurations have been inserted. This is ~~the~~an aim of the last two steps:

- ~~—~~—RISC Code Generation and
- 10 • ~~—~~—RISC Code Scheduling.

The first one ~~produces~~may produce code for the host processor and the second one ~~optimizes~~may optimize it further by looking for a better scheduling using software pipelining for instance.

15

4.3—XPP Compiler for Loops

~~Figure 52 describes the~~Fig. 25 illustrates a detailed architecture and an internal processing of the XPP Compiler. It is a complex cooperation between program transformations, included in the XPP Loop optimizations, a temporal partitioning phase, NML code generation and the
20 mapping of the configuration on the PACT XPP.

~~Figure 52: Detailed Architecture of the XPP Compiler~~

First, loop optimizations targeted at the PACT XPP ~~are~~may be applied to try to produce innermost loop bodies that can be executed on the array of processors. If this is the case, the
25 NML code generation phase ~~is~~may be called; ~~if~~if. If not, then temporal partitioning ~~is~~may be applied to get several configurations for the same loop. After NML code generation and the mapping phase, it can also happen that a configuration will not fit on tike PACT XPP. In this case, the loop optimizations ~~are~~may be applied again with respect to the reasons of failure of the NML code generation or of the mapping. If this new application of loop optimizations
30 does not change the code, temporal partitioning ~~is~~may be applied; ~~Furthermore we keep track of~~if the number of attempts for the NML Code Generation and the mapping, ~~if~~may be kept track of. If too many attempts are made, and ~~we still do not obtain~~is still not obtained, ~~we break~~may be broken the process, and the loop nest ~~will~~may be executed by the host processor.

4.3.1—Temporal Partitioning

Temporal partitioning ~~splits~~may split the code generated for the PACT XPP ~~into~~ several configurations if the number of operations, *i. e.*, the size of the configuration, to be executed in a loop nest exceeds the number of operations executable in a single configuration. This transformation is called loop dissevering [6]. ~~These configurations are.~~ See, for example, João M.P. Cardoso & Markus Weinhardt, “XPP-VC: A C Compiler with Temporal Partitioning for the PACT-XPP Architecture,” *Proceedings of the 12th International Conference on Field-Programmable Logic and Applications, FPL’2002*, 2438 LNCS, 864-874 (2002). These configurations may be then integrated in a loop of configurations whose number of execution corresponds to the iteration range of the original loop.

4.3.2—Generation of NML Code

~~This step takes~~Generation of NML code may take as input an intermediate form of the code produced by the XPP Loop optimizations step, together with a dataflow graph built upon it. NML code can then be produced by using tree- or DAG-pattern matching techniques.

4.3.3—Mapping Step

~~This~~A mapping step takes~~may take~~ care of mapping the NML₇ modules on the PACT XPP by placing the operations on the ALUs, FREGs, and BREGs, and routing the data through the buses.

4.4—XPP Loop Optimizations Driver

~~The~~A goal of loop optimizations used for the PACT XPP ~~are now described.~~ Their goal is to extract as much parallelism as possible from the loop nests in order to execute them on the PACT XPP by exploiting the ALU-PAEs as effectively as possible and to avoid memory bottlenecks with the IRAMs. The following sections explain how they ~~are~~may be organized and how to take into account the architecture for applying the optimizations.

4.4.1—Organization of the System

~~Figure 53 below presents the organization~~Fig. 26 provides a detailed view of the XPP loop optimizations, including their organization. The transformations ~~are~~may be divided in six groups. Other standard optimizations and analysis ~~are~~may be applied in-between. Each

group could be called several times. Loops over several groups can also occur if needed. The number of iterations for each driver loop can be of constant value or determined at compile time by the optimizations ~~itself~~themselves, (e.g., repeat until a certain code quality is reached). In the first iteration of the loop, it can be checked if loop nests are usable for the PACT XPP, ~~it~~. It is mainly directed to check the loop bounds etc. For instance, if the loop nest is well-formed and the data dependence graph does not prevent optimization, but the loop bounds are unknown, then, in the first iteration loop, tiling is~~may be~~ applied to get an innermost that is easier to handle and can be better optimized, and in the second iteration, loop normalization, if conversion, loop interchange and other optimizations can be, applied to effectively optimize the inner-most loops for the PACT XPP. Nevertheless, this has not been necessary until now with the examples presented ~~in the next chapters~~. below.

With reference to Fig. 26, Group I ~~ensures~~may ensure that no procedure calls occur in the loop nest. Group II ~~prepares~~may prepare the loop bodies by removing loop-invariant instructions and conditional instruction to ease the analysis. Group III ~~generates~~may generate loop nests suitable for the data dependence analysis. Group IV ~~contains~~may contain optimizations to transform the loop nests to get data dependence graphs that are suitable for vectorization. Group V ~~contains~~may contain optimizations that ensure that the innermost loops can be executed on the PACT XPP. Group VI ~~contains~~may contain optimizations that further extract parallelism from the loop bodies. Group VII ~~contains~~may contain optimizations more towards optimizing the usage of the hardware itself.

In each group, the application of the optimizations ~~depends~~may depend on the result of the analysis and the characteristics of the loop nest. For instance, it is clear that not all transformations in Group IV are applied. It depends on the data dependence graph computed before.

Figure 53: Detailed View of the XPP Loop Optimizations

4.4.2—Loop Preparation

The optimizations of Groups I, II and III of the XPP compiler may generate loop bodies without procedure calls, conditional instructions and induction variables other than loop control variables. Thus, loop nests, where the innermost loops are suitable for execution on

the PACT XPP, ~~are~~may be obtained. The iteration ranges ~~are~~may be normalized to ease data dependence analysis and the application of other code transformations.

4.4.3—Transformation of the Data Dependence Graph

- 5 The optimizations of Group IV ~~are~~may be performed to obtain innermost loops suitable for vectorization with respect to the data dependence graph. Nevertheless, a difference with usual vectorization is that a dependence cycle, ~~that~~which would normally prevent any vectorization of the code, does not prevent the optimization of a loop nest for the PACT XPP. If a cycle is due to an anti-dependence, then it could be that it ~~won't~~will not prevent
- 10 optimization of the code as stated in [7]. ~~Furthermore~~Markus Weinhardt & Wayne Luk, "Pipeline Vectorization," IEEE Transactions on Computer-Aided Design of integrated Circuits and Systems, 20(2):234-248 (February 2001). Furthermore, dependence cycles will not pre-vent vectorization for the PACT XPP- when it consists only of a loop-carried true dependence on the same expression. If cycles with distance k occur in the data dependence
- 15 graph, then this ~~an~~can be handled by holding k values in registers. This optimization is of the same class as cycle shrinking.

- Nevertheless, limitations due to the dependence graph exist. Loop nests cannot be handled if some dependence distances are not constant, or unknown. If only a few
- 20 ~~dependences~~dependencies prevent the optimization of the whole loop nest, this could be overcome, by using the traditional vectorization algorithm that sorts topologically the strongly connected components of the data dependence graph (statement reordering), and then ~~apply~~applying loop distribution. This way, loop nests, which can be handled by the PACT XPP and some by the host processor, can be obtained.

25

4.4.4—Influence of the Architectural Parameters

- Some hardware specific parameters may influence the application of the loop transformations. The number of operations and memory accesses, that a loop body performs, is may be estimated at each step. These parameters may influence loop unrolling, strip-
- 30 mining, loop tiling and also loop interchange (iteration range).

The table below lists the, parameters that may influence the application of the optimizations. For each of them, two data are given: a starting value computed from the loop, and a restriction value which is the value the parameter should reach or should not exceed after the

application of the optimizations. Vector length depicts the range of the innermost loops, *i.e.*, the number of elements of an array accessed in the loop body. Reused data set size represents the amount of data that must fit in the cache. I/O IRAMs, ALU, FREG, BREG stand for the number of IRAMs, ALUs, FREGs, and BREGs, respectively ~~that constitute~~ of the PACT XPP. The dataflow graph width represents the number of operations that can be executed in parallel in the same pipeline stage. The dataflow graph height represents the length of the pipeline. Configuration cycles amounts to the length of the pipeline, and to the number of cycles dedicated to the control. The application of each optimization may

- —decrease a parameter's value (-),
- —increase a parameter's value (+),
- —not influence a parameter (id), or
- —adapt a parameter's value to fit into the goal size (make fit).

Furthermore, some resources must be kept for control in the configuration; ~~this~~. This means that the optimizations should not make the needs exceed more than 70-80% each resource.

Parameter	Goal	Starting Value
Vector length	IRAM size (256 words)	Loop count
Reused data set size	Approx. cache size	Algorithm analysis/loop sizes
I/O IRAMs	PACT size (16)	Algorithm inputs + outputs
ALU	PACT size (< 64)	ALU opcode estimate
BREG	PACT size (< 80)	BREG opcode estimate
FREG	PACT size (< 80)	FREG opcode estimate
Data flow graph width	High	Algorithm data flow graph
Data flow graph height	Small	Algorithm data flow graph
Configuration cycles	≤ command line parameter	Algorithm analysis

~~Here are some additional~~ Additional notations used² in the following descriptions are as follows. ~~Let n be~~ is the total number of processing elements available, r , is the width of the dataflow graph, in , is the maximum number of input values in a cycle, and out , is the maximum number of output values possible in a cycle. On the PACT XPP, n is the number of ALUs, FREGs and BREGs available for a configuration, r is the number of ALUs, FREGs and BREGs that can be started in parallel in the same pipeline stage, and, in and out amount to the number of available IRAMs. As IRAMs have 1 input port and 1 output port, the number of IRAMs yields directly the number of input and output data.

The number of operations of a loop body is may be computed by adding all logic and arithmetic operations occurring in the instructions. The number of input values is the number of operands of the instructions regardless of address operations. The number of output values is the number of output operands of the instructions regardless of address operations. To

5 determine the number of parallel operations, input and output values, and the dataflow graph must be considered. The effects of each transformation on the architectural parameters are now presented in detail.

Loop Interchange

- 10 Loop interchange is may applied when the innermost loop has a too narrow iteration range. In that case, loop interchange ~~allows to have~~ may allow for an innermost loop with a ~~in ore~~ more profitable iteration range. It can also be influenced by— the layout of the data in memory. It can be profitable to data locality to interchange two loops to get a more practical way to access arrays in the cache and therefore prevent cache misses. It is of course also influenced
- 15 by data ~~dependences~~ dependencies as explained ~~earlier.~~ above.

Parameter	Effect
Vector length	+
Reused data set size	make fit
I/O IRAMs	id
ALU	id
BREG	id
FREG	id
Data flow graph width	id
Data flow graph height	id
Configuration cycles	-

Loop Distribution

- Loop distribution is may be applied if a loop body is too big to fit on the PACT XPP. Its A main effect of loop distribution is to reduce the processing elements needed by the
- 20 configuration. Reducing the need for IRAMs can only be a side effect.

Parameter	Effect
Vector length	id
Reused data set size	id
I/O IRAMs	make fit
ALU	make fit
BREG	make fit
FREG	make fit
Data flow graph width	-
Data flow graph height	-

Configuration cycles	-
----------------------	---

Loop Collapsing

Loop collapsing can be used to make the loop body use more memory resources. As several dimensions are merged, the iteration range is increased and the memory needed ~~is~~ increased as well.

Parameter	Effect
Vector length	+
Reused data set size	+
I/O IRAMs	+
ALU	id
BREG	id
FREG	id
Data flow graph width	+
Data flow graph height	+
Configuration cycles	+

Loop Tiling

Loop tiling, as multi-dimensional strip-mining, is influenced by all parameters, ~~it is~~. It may be especially useful when the iteration space is by far too big to fit in the IRAM, or to guarantee maximum execution time when the iteration space is unbounded (~~see Section 4.4.6~~). See the discussion below under the heading “Limiting the Execution Time of a Configuration.” It can then make the loop body fit with respect to the resources of the PACT XPP, namely the IRAM and cache line sizes. The size of the tiles for strip-mining and loop tiling can be computed ~~like this~~:

$$\text{tile size} = \text{resources available for the loop body} / \text{resources necessary for the loop body.}$$

The resources available for the loop body are the whole resources of the PACT XPP for this configuration. A tile size can be computed for the data and another one for the processing elements, ~~the~~. The final tile size is then the minimum between these two. For instance, when the amount of data accessed is larger than the capacity of the cache, loop tiling ~~can~~ may be applied ~~like below~~. according to the following example code for loop tiling for the PACT XPP.

```

for (i=0; i<=1048576; i++) for (i=0; i<=1048576; i+= CACHE_SIZE)
    <loop body>                for (j=0; j<CACHE_SIZE; j+=IRAM_SIZE)
                                for (k=0; k<IRAM_SIZE; k++)

```


<tilted loop body>

Figure 54: Example of loop tiling for the PACT XPP

Parameter	Effect
Vector length	make fit
Reused data set size	make fit
I/O IRAMs	id.
ALU	id
BREG	id
FREG	id
Data flow graph width	+
Data flow graph height	+
Configuration cycles	+

5 Strip-Mining

Strip-mining is may be used to make the amount of memory accesses of the innermost loop fit with the IRAMs capacity. The processing elements do not usually represent a problem as the PACT XPP has 64 ALU-PAEs which should be sufficient to execute any single loop body. Nevertheless, the number of operations can be also taken into account the same way as the data.

Parameter	Effect
Vector length-	make fit
Reused data set size	id
I/O IRAMs	id
ALU	id
BREG	id
FREG	id
Data flow graph width	+
Data flow graph height	-
Configuration cycles	-

Loop Fusion

Loop fusion is may be applied when a loop body does not use enough resources. In this case, several loop bodies can be merged to obtain a configuration using a larger part of the available resources.

Parameter	Effect
Vector length	id
Reused data set size	id
I/O IRAMs	+-
ALU	+
BREG	+

FREG	+
Data flow graph width	id
Data flow graph height :	+
Configuration cycles	+

Scalar Replacement

The amount of memory needed by the loop body should always fit in the IRAMs.

- 5 ~~Thanks~~Due to this a scalar replacement optimization, some input or output data represented by array references, that should be stored in IRAMs, ~~are~~ may be replaced by scalars, that are either stored in FREGs or kept on buses.

Parameter	Effect
Vector length	+
Reused data set size	id
I/O IRAMs	id
ALU	id
BREG	id
FREG	id
Data flow graph width	+
Data flow graph height	-
Configuration cycles	id

Loop Unrolling

- 10 Loop unrolling, loop collapsing, loop fusion and loop distribution ~~are~~ may be influenced by the number of operations of the body of the, loop nest² and the number of data inputs and outputs of these operations, as they modify the size of the loop body. The number of operations should always be smaller than n , and the number of input and output data should always be smaller than in and out .

Parameter	Effect
Vector length	id
Reused data set size	id,
I/O IRAMs	+
ALU	+
BREG	+
FREG	+
Data flow graph width	id
Data flow graph height	+
Configuration cycles	+

15

Unroll-and-Jam

Unroll-and-Jam ~~consists in~~may include unrolling an outer loop and then merging the inner loops. It must compute the unrolling degree u with respect to the number of input memory accesses m and output memory accesses p in the inner loop. The following inequality must

- 5 hold: $u * m \leq in \wedge u * p \leq out$. Moreover, the number of operations of the new inner loop must also fit on the PACT XPP.

Parameter	Effect
Vector length	id
Reused data set size	+
I/O IRAMs	+
ALU	+
BREG	+
FREG	+
Data flow graph width	id
Data flow graph height	+
Configuration cycles	+

4.4.5—Optimizations Towards Hardware Improvements

- 10 At this step other optimizations, specific to the PACT XPP, can be made. These optimizations deal mostly with memory problems and dataflow considerations. This is the case of shift register synthesis, input data duplication (similar to scalar privatization), or loop pipelining.

15 Shift Register Synthesis

- This A shift register synthesis optimization deals with array accesses that occur during the execution of a loop body. When several values of an array are alive for different iterations, it can be convenient to store them in registers, rather than accessing memory each time they are needed. As the same value must be stored in different registers depending on the number of
- 20 iterations it is alive, a value shares several registers and flows from a register to another at each iteration. It is similar to a vector register allocated to an array access with the same value for each element. This optimization is performed directly on the dataflow graph by inserting nodes representing registers when a value must be stored in a register. In the PACT XPP, it amounts to ~~store~~storing it in a data register. A detailed explanation can be found in
- 25 [1]—Markus Weinhardt & Wayne Luk, “Memory Access Optimization for Reconfigurable Systems,” IEEE Proceedings Computers and Digital Techniques, 48(3) (May 2001).

Shift register synthesis ~~is~~may be mainly suitable for small to medium amounts of iterations where values are alive. Since the pipeline length increases with each iteration for which the value has to be buffered, the following method is better suited for medium to large distances between accesses in one input array.

5

Nevertheless, this method ~~works~~may work very well for image processing algorithms which mostly alter a pixel by analyzing itself and its surrounding neighbors.

Parameter	Effect
Vector length	+
Reused data set size	id
I/O IRAMs	id
ALU	id
BREG	id
FREG	id
Data flow graph width	+
Data flow graph height	-
Configuration cycles	id

Input Data Duplication

- 10 This An input data duplication optimization is orthogonal to shift register synthesis. If different elements of the same array are needed concurrently, instead of storing the values in registers, the same values ~~are~~may be copied in different IRAMs. The advantage against shift register synthesis is the shorter pipeline length, and therefore the increased parallelism, and the unrestricted applicability. On the other hand, the cache-IRAMM bottle-neck can affect
- 15 the performance of this solution, depending on the amounts of data to be moved. Nevertheless ~~we assume~~, it is assumed that cache IRAM transfers are negligible to transfers in the rest of the memory hierarchy.

Parameter	Effect
Vector length	+
Reused data set size	id
I/O IRAMs	id
ALU	id
BREG	id
FREG	id
Data flow graph width	+
Data flow graph height	-
Configuration cycles	id

Loop Pipelining

A loop optimization ~~consists in~~ pipelining optimization may include synchronizing operations by inserting delays in the dataflow graph. These delays ~~are~~ may be registers. For the PACT XPP, it amounts to ~~store~~ storing values in data registers to delay the operation using them. This is the same as pipeline balancing performed by xmap.

Parameter	Effect
Vector length	+
Reused data set size	id
I/O IRAMs	id
ALU	id
BREG	id
FREG	id
Data flow graph width	+
Data flow graph height	-
Configuration cycles	+

Tree Balancing

A tree balancing optimization ~~consists in~~ may include balancing the tree representing the loop body. It ~~reduces~~ may reduce the depth of the pipeline, thus reducing the execution time of an iteration, and ~~increases~~ may increase parallelism.

Parameter	Effect
Vector length	+
Reused data set size	id
I/O IRAMs	id
ALU	id
BREG	id
FREG	id
Data flow graph width	+
Data flow graph height	-
Configuration cycles	-

4.4.6—Limiting the Execution Time of a Configuration

The execution time of a configuration must be controlled. This is ensured in the compiler by strip-mining and loop tiling that take care that not more input data ~~as~~ than the IRAMs' capacity come in the PACT XPP in a cycle. This way the iteration-range of the innermost loop that is— executed on the PACT XPP is limited, and therefore its execution time.

Moreover, partitioning ensures that loops, whose execution count can be computed at run time, are going to be executed on the PACT XPP. This condition is trivial for for-loops, but for while-loops, where the execution count cannot ~~de~~ be determined statically, a

transformation like ~~sketched~~ exemplified by the code below can be applied. As a result, the inner for-loop can be handled by the PACT XPP.

```

        while (ok) {
            <loop body>
5      i++) {
            }
        }
                                while (ok)
                                for (i=0; i<100 && ok;
                                <loop body>

```

Figure 55: Transformation of while-loops

10 ~~5~~ CASE STUDIES

~~5.1~~ 3x3 Edge Detector

~~5.1.1~~ Original Code

~~Source Code:~~

The following is source code:

```

15  #define VERLEN 16
    #define HORLEN 16
    main() {
        int v, h, inp;
        int p1[VERLEN][HORLEN];
20  int p2[VERLEN][HORLEN];
        int htmp, vtmp, sum;

        for (v=0; v<VERLEN; v++)      // loop nest 1
            for (h=0; h<HORLEN; h++) {
25  scanf("%d", &p1[v][h]); // read input pixels to p1
                p2[v][h] = 0;          // initialize p2
            }

        for (v=0; v<=VERLEN-3; v++) { // loop nest 2
30  for (h=0; h<=HORLEN-3; h++) {
            htmp = (p1[v+2][h] - p1[v][h]) +
                    (p1[v+2][h+2] - p1[v][h+2]) +
                    2 * (p1[v+2][h+1] - p1[v][h+1]);

```

```

    if (htmp < 0)
        htmp = -htmp ;

    vtmp = (p1[v][h+2] - p1[v][h]) +
5        (p1[v+2][h+2] - p1[v+2][h]) +
        2 * (p1[v+1][h+2] - p1[v+1][h]) ;
    if (vtmp < 0)
        vtmp = -vtmp ;

10    sum = htmp + vtmp;
    if (sum > 255)
        sum = 255;
    p2[v+1][h+1] = sum;
}
15 }
    for (v=0; v<VERLEN; v++)          // loop nest 3
        for (h=0 ; h<HORLEN; h++)
            printf("%d\n", p2[v][h]) ; // print output pixels from p2
}

```

5.1.2—Preliminary Transformations

Interprocedural Optimizations

The first step normally invokes interprocedural transformations like function dining and loop pushing. Since no procedure calls are within the loop body, these transformations are not applied to this example.

Partitioning

The partitioning algorithm chooses which code runs on the RISC processor and which code runs on the XPP. Since ~~we only consider~~ inner loops are considered to run on the XPP, the basic blocks are annotated with the loop nest depth. Thus, basic blocks which are not in a loop are separated out. Furthermore, function calls within a loop body prevent a loop to be considered for running on the XPP.

In our benchmark, the loop nests 1 and 3 are marked as to run on the RISC host because of the function call. In the following sections they are not considered any further.

- 5 It is to say that at this compilation stage it is not predictable if the remaining loop nests can be synthesized for the XPP. ~~We just separated~~Just the ones which definitely cannot run on it; ~~others were separated.~~ Others may follow, since running the code on the RISC CPU is always the reassurance in ~~our~~this strategy.

10 **Loop Analysis and Normalization**

The code upon has already normalized loops. Nevertheless, it is more likely that human written code would ~~look like~~be approximately as follows:

```

for (v=1; v < VERLEN - 1; v++) {
    for (h=1; h < HORLEN - 1; h++) {
15         htmp = (p1[v+1][h-1] - p1[v-1][h-1]) +
                (p1[v+1][h+1] - p1[v-1][h+1]) +
                2 * (p1[v+1][h] - p1[v-1][h]);
        if (htmp < 0)
            htmp = -htmp;
20
        vtmp = (p1[v-1][h+1] - p1[v-1][h-1]) +
                (p1[v+1][h+1] - p1[v+1][h-1]) +
                2 * (p1[v][h+1] - p1[v][h-1]);
        if (vtmp < 0)
25            vtmp = -vtmp;

        sum = htmp + vtmp;
        if (sum > 255)
            sum = 255;
30        p2[v+1][h+1] = sum; }
    }
}
```


Although seen at first sight by a human reader, it is not obvious for the compiler that the loop is well formed. Therefore ~~it is tried to normalize~~, normalizing of the loop is attempted.

If the original loop induction variable is called i with the increment value s and lower and upper loop bounds l and u , respectively, then the normalized loop with the induction variable i'_l and the upper bound u'_l (the lower bound l'_l is 0 -by- definition) is transformed as follows:

- ~~■~~—The upper bound calculates to $u'_l = (u-l)/s$.
- ~~■~~—All occurrences of i are replaced by $l + i'_l * s$.

Applied to the code above, the loop statement for ($v=1$; $v < \text{VERLEN} - 1$; $v++$) with the lower bound $vl = 1$, the upper bound $vu = 14$ (< 15 means ≤ 14 in integer arithmetic) and the increment $vs = 1$ transforms to

```
for (vn=0; vn <= (vu - vl) / vs; vn++)
```

or simplified

```
for (vn= 0 ; vn <= 13; vn++ )
```

The 'h-loop' is transformed equally, issuing the original code.

Idiom Recognition

In the second step, idiom recognition finds the `abs()` and `min()` structures in the loop body. ~~Please note~~Note that although the XPP has no `abs` opcode, it can easily be synthesized and should therefore be produced to simplify the internal representation ~~(otherwise, (Otherwise,~~ if-conversion has to handle this case which increases the complexity).)

Therefore, the code after idiom recognition ~~looks like~~ is approximately as follows (`abs()` and `min()` are compiler known functions which are directly mapped to XPP opcodes or predefined NML modules):

```
for (v=0; v<=16-3; v++) {
  for (h=0; h<=16--3; h++) {
    htmp = (p1 [v+2] [h] - p1 [v] [h]) +
            (p1 [v+2] [h+2] - p1 [v] [h+2]) +
            2 * (p1 [v+2] [h+1] - p1 [v] [h+1]) ;
    htmp = abs (htmp) ;
```

```

    vtmp = (p1[v][h+2] - p1[v][h]) +
            (p1[v+2][h+2] - p1[v+2][h]) +
            2 * (p1[v+1][h+2] - p1[v+1][h]);
5    vtmp = abs(vtmp);

    sum = min(htmp + vtmp, 255);
    p2[v+1][h+1] = sum;
}
10 }
```

Dependency Analysis

```

for(v=0; v<=16-3; v++){
    for(h=0; h<=16-3; h++){
15
S1        htmp = (p1[v+2][h] - p1[v][h]) +
                (p1[v+2][h+2] - p1[v][h+2]) +
                2 * (p1[v+2][h+1] - p1[v][h+1]);
S2        htmp = abs(htmp);
20
S3        vtmp = (p1[v][h+2] - p1[v][h]) +
                (p1[v+2][h+2] - p1[v+2][h]) +
                2 * (p1[v+1][h+2] - p1[v+1][h]);
S4        vtmp = abs(vtmp);
25
S5        sum = min(htmp + vtmp, 255);

S6        p2[v+1][h+1] = sum;
    }
30 }
```

~~Figure 56 The expression tree of the edge 3x3 inner loop body~~

There are no loop carried dependencies which prevent pipeline vectorization. The loop independent scalar dependencies do not prevent pipeline vectorization since the transformation does not disturb the order of reads and writes. Furthermore, forward expression substitution / dead code elimination will remove the scalars completely.

5.1.3—Pre Code Generation Transformations

Forward Expression Substitution / Dead Code Elimination

The lack of uses of `htmp`, `vtmp` and `sum` after the loop nest allows forward expression substitution along with dead code elimination to place the whole calculation into one statement.

```
p2 [v+1] [h+1] = min (abs ((p1 [v+2] [h1] - p1 [v] [h]) +
                        (p1 [v+2] [h+2] - p1 [v] [h+2])) +
                    2 * (p1 [v+2] [h+1] - p1 [v] [h+1])) +
                    abs ((p1 [v] [h+2] - p1 [v] [h]) +
                        (p1 [v+2] [h+2] - p1 [v+2] [h]) +
                        2 * (p1 [v+1] [h+2] - p1 [v+1] [h])), 255);
```

The scalar accesses then disappear completely.

Mapping to IRAMs

The array accesses are mapped to IRAMs. At this stage the IRAM numbers are chosen arbitrarily, ~~the~~. The actual mapping to XPP IRAMs is done later.

Therefore ~~we rename~~ `p1`, `p1[v+x][h+y]` and `p2[v+X][h+y]` are renamed to `iramN[y]`, (e.g., `p1[v+2][h]` to `iram2[0]`). ~~The~~ Accordingly, the code reads then is

```
iram3 [ 1 ] = min (abs (iram2 [ 0 ] - iram0 [0]) +
                  (iram2 [2] - iram0 [2]) +
                  2 * (iram2 [1] - iram0 [1]) +
                  abs (iram0 [2] - iram0 [0]) +
                  (iram2 [2] - iram2 [0]) +
                  2 * (iram1 [2] - iram1 [0]), 255);
```

Tree Balancing

Fig. 27 shows an expression tree of an edge 3x3 inner loop body. The visualized expression tree in Figure 56 of Fig. 27 shows another valuable optimization before matching the tree.

Since the depth of the tree determines the length of the synthesized pipeline, another

5 simplification can decrease this depth. In both of the main sub trees, the operands of the commutative add expressions can be interchanged to reduce the overall tree depth. Figure 57
OneA resulting expression tree is shown in Fig. 28. In Fig. 28, one of the sub trees is shown
before and after balancing. The numbers represent the annotated maximum ~~tree~~ tree depth from the node to its deepest child leaf node.

10 The resulting expression tree is shown in Figure 57.

5.1.4—XPP Code generation

Pipeline Synthesis

As already stated, the pipeline is synthesized by a dynamic programming tree matcher. In
15 contrast to sequential processors, it does not generate instructions and register references, but PAE opcodes and port connections. The Fig. 29 shows the main calculation network is shown
in Figure 58 of the edge 3x3 configuration. The MULTI-SORT combination does the abs() calculation, while the SORT does the min() calculation. The input data preparation network is not shown in this figure. The Fig. 29. Fig. 30 shows the case of synthesized shift registers
20 are shown in Figure 59, while the variant with duplicated input data simply consists
of includes an IRAM for each input channel in Figure 58-Fig. 29. With respect to Fig. 30,
there is one input after the shift register synthesis. The leftmost input contains p1[][h], the
middle one contains p1[][h+1], and the rightmost one contains p1[][h+2].

25 Although this is straight forward, there remains the question how to access the different offsets of the vector register accesses. Although the RAM-PAEs are dual ported, it is obvious that it is not possible to read different addresses concurrently.

Since it is not efficient to synthesize a configuration which generates the different addresses
30 sequentially and demultiplexes the read operands into different branches of the data flow, other arrangements have to be made.

The two possibilities to access input data ~~presented in subsection 4.4.5~~ discussed above under the heading “Optimizations Towards Hardware Improvements” yield the following in RISC pseudo code and XPP utilization. ~~—The pseudo code running on the RISC core looks like is approximately:~~

```

5  XPPPreload (config)
    for (v=0; v<=16-3; v++) {
        XPPPreload(0, &p1[v], 16)

        XPPPreload(1, &p1[v+1], 16)

        XPPPreload(2, &p1[v+2], 16)
10    XPPPreloadClean (3, &p2[v+1], 16)

        XPPEecute (config, IRAM(0), IRAM(1), IRAM(2), IRAM(3))
    }

```

for shift register synthesis and approximately:

```

15 XPPPreload(config)
    for(v=0; v<=16-3; v++){
        XPPPreload(0, &p1[v], 16)
        XPPPreload(1, &p1[v], 16)
        XPPPreload(2, &p1[v], 16)
20    XPPPreload(3, &p1[v+1], 16)
        XPPPreload(4, &p1[v+1], 16)
        XPPPreload(5, &p1[v+2], 16)
        XPPPreload(6, &p1[v+2], 16)
        XPPPreload(7, &p1[v+2], 16)
25    XPPPreloadClean(3, &p2[v+1], 16)
        XPPEecute(config, IRAM(0), IRAM(1), IRAM(2), IRAM(3),
                    IRAM(4), IRAM(5), IRAM(6), IRAM(7))
    }

```

~~for data duplication, respectively.~~

30

~~Figure 58 The main calculation network of the edge3x3 configuration. The MULT SORT combination does the abs() calculation while the SORT does the mind calculation.~~

~~Figure 59 One input after shift register synthesis. The leftmost input contains $p1[] [h]$, the middle one $p1[] [h+1]$ and the rightmost $p1[] [h+2]$, respectively.~~

The values for place & route and simulation are compared in the following table. Note that a common RISC DSP with two MAC units and hardware loop support needs about 4000 cycles for the same code. This comparison does not account for cache misses. Furthermore, it is obvious, that the number of input values is very small in this example and the DSP calculation time is proportional to that number. The XPP performance on the other hand will improve with the number of input values. Therefore, the XPP performance will be more impressive with bigger image sizes.

Parameter	Value (shift register synthesis)		Value (data duplication)	
Vector length	16		16	
Reused data set size	256		256	
I/O IRAMs	3 I + 1 O = 4		8 I + 1 O = 9	
ALU	27		21	
BREG	21 (1 defined + 20 route)		10 (1 defined + 9 route)	
FREDFREg	22 (9 defined + 23 route)		19 (3 defined + 16 route)	
Data flow graph width	14		14	
Data flow graph height	3 (shift registers) + 8 (calculation)		8 (calculation)	
Configuration cycles (simulated)	configuration	2262	configuration	2145
	preloads [†] (assuming 4 words/cycle burst transfer)	14*3*4 168	preloads	8*8*4 256
	cycles	14*57 798	cycles	14*52 728
	sum	3228	sum	3129

5.1.5 Enhancing Parallelism

After the synthesis, the configuration calculating the inner loop utilizes 27 ALUs and 4 IRAMs for shift register synthesis and 21 ALUs and 9 IRAMs for data duplication,

[†] assuming 4 words/cycle burst transfer

respectively. Assuming ~~aan~~ XPP64 core, this leaves plenty of room for further optimizations. Nevertheless, since all optimizations enhancing parallelism are performed before the synthesis takes place, it is crucial that they estimate the needed resources and the benefit of the transformation very carefully. Furthermore, they have to account for both input
5 preparation strategies to estimate correct values.

Loop Unrolling

Fully unrolling the inner loop would not lead to satisfying results, because the number of inputs and outputs increases dramatically. ~~That~~This means data duplication would not be
10 applicable and shift register synthesis would exhaust most of the benefits of the parallelism by producing a very long pipeline for each data flow graph. Although partial unrolling of the inner ~~to- ϕ -loop~~ would be applicable, it promises not much benefit for the area penalty introduced.

15 Loop unrolling the outer loop is also not applicable since it produces a further configuration. Nevertheless, a related transformation could do a good job on this loop nest.

Unroll-and-Jam

The unroll-and-jam algorithm enhances parallelism and also improves IRAM usage. It brings
20 pairs of iterations together ideally reusing IRAM outputs and calculation results. The algorithm partially unrolls the outer loop and fuses the originated inner loops. Before the unroll-and-jam is performed, the so-called unroll-and-jam factor must be determined, which denominates the unrolling factor of the outer loop. This is mainly influenced by the number of ALUs n ($= 64$ assuming XPP64) and calculates to

$$25 \quad c_{\text{unroll-and-jam}} = \frac{n_{\text{XPP}}}{n_{\text{inner loop}}} = \frac{64}{27} = 2 \text{ (integer division)}$$

Thus the source code would be transformed to:—;

```
for (v=0; v<=VERLEN-3; v+=2) {
    for (h=0; h<=HORLEN-3; h++) {
30         p2[v+1][h+1] = min(abs((p1[v+2][h] - p1[v][h]) +
                                (p1[v+2][h+2] - p1[v][h+2]) +
                                2 * (p1[v+2][h+1] - p1[v][h+1])) +
                                abs((p1[v][h+2] - p1[v][h]) +
```

```

                    (p1[v+2][h+2] - p1[v+2][h]) +
                    2 * (p1[v+1][h+2] - p1[v+1][h])), 255);
p2[v+2][h+1] = min(abs((p1[v+3][h] - p1[v+1][h]) +
                    (p1[v+3][h+2] - p1[v+1][h+2]) +
5                2 * (p1[v+3][h+1] - p1[v+1][h+1])) +
                    abs((p1[v+1][h+2] - p1[v+1][h]) +
                    (p1[v+3][h+2] - p1[v+3][h]) +
                    2 * (p1[v+2][h+2] - p1[v+2][h])), 255);
    }
10 }

```

The transformation introduces additional accesses to $p1[v+3][h]$, $p1[v+3][h+2]$, $p1[v+3][h+1]$, and $p1[v+1][h+1]$ (the former hole in the access pattern) as well as a write access to $p2[v+2][h+1]$. ~~That~~This means 2 IRAMs more for shift register synthesis (one

15 input, one output) and 5 IRAMs more for data; duplication (4 input, 1 output), while performance is doubled.

Parameter	Value (shift register synthesis)		Value (data duplication – no IRAM placement)	
Vector length	16		16	
Reused data set size	256		256	
I/O, IRAMs	4 I + 2 O = 6		12 I + 2 O = 14	
ALU	45		37	
BREG	31 (12 defined + 19 route)		42 (4 defined + 38 route)	
FREG	29 (1 defined + 28 route)		18 (1 defined + 17 route)	
Data flow graph width	14		14	
Data flow graph height	3 (shift registers) + 8 (calculation)		8 (calculation)	
Configuration cycles (simulated)	configuration	2753	configuration	2754
	preloads	7*4*4 112	preloads	7*12*4
	cycles	7*53 371	cycles	336
	sum	3236	sum	7*69 483
				3573

Parameter	Value (data duplications – with IRAM placement)		
Vector length	16		
Reused data set size	256		
I/O IRAMs	12 I + 2 O = 14		
ALU	37		
BREG	36 (4 defined + 32 route)		
FREG	24 (1 defined + 23 route)		
Data flow graph width	14		
Data flow graph height	3 (shift registers) + 8 (calculation)		
Configuration cycles (simulated)	configuration	2768	
	preloads	7*12*4	336
	cycles	7*51	357
	sum	3461	

The simulated results are shown in the table above. ~~Please note~~ Note the differences of the two columns labeled with “data duplication.”: The first used xmap to place the IRAMs, while in the second, the IRAMs were placed by hand using a greedy algorithm which places IRAMs that are operands of the same operator in one line (as long as this is possible). The second solution improved the iteration cycles by 18. This shows that IRAM placement has a great impact to the final performance.

The traditional unroll-and-jam algorithm uses loop peeling to split the outer loop in a preloop and an unroll-able main loop to handle odd loop counts. When ~~we assume~~, for instance, $n = 128$ is assumed, the unroll-and-jam factor would calculate to $c_{unroll-and-jam} = \frac{128}{27} = 4$.

Since the outer loop count (14) is not a multiple of 4, the algorithm virtually peels off the first two iterations and fuses the two loops at the end adding guards to the inner loop body. Then the code looks ~~like~~ approximately as follows (guards *emphasized*):

for

```
for (v=0; v<=VERLEN-5; v+=4) {
```

```
    for (h=0; h<=HORLEN-3; h++) {
```

```
        p2[v+1][h+1] = min(abs((p1[v+2][h] - p1[v][h]) +
```

```

                    (p1[v+2][h+2] - p1[v][h+2]) +
                    2 * (p1[v+2][h+1] - p1[v][h+1])) +
                    abs((p1[v][h+2] - p1[v][h]) +
                    (p1[v+2][h+2] - p1[v+2][h]) +
5                    2 * (p1[v+1][h+2] - p1[v+1][h])), 255);
    p2[v+2][h+1] = min(abs((p1[v+3][h] - p1[v+1][h]) +
                    (p1[v+3][h+2] - p1[v+1][h+2]) +
                    2 * (p1[v+3][h+1] - p1[v+1][h+1])) +
                    abs((p1[v+1][h+2] - p1[v+1][h]) +
10                    (p1[v+3][h+2] - p1[v+3][h]) +
                    2 * (p1[v+2][h+2] - p1[v+2][h])), 255);
    if(v>0) p2[v+3][h+1] = min(abs((p1[v+4][h] - p1[v+2][h]) +
                    (p1[v+4][h+2] - p1[v+2][h+2]) +
                    2 * (p1[v+4][h+1] - p1[v+2][h+1])) +
15                    abs((p1[v+2][h+2] - p1[v+2][h]) +
                    (p1[v+4][h+2] - p1[v+4][h]) +
                    2 * (p1[v+3][h+2] - p1[v+3][h])), 255);
    if(v>1) p2[v+4][h+1] = min(abs((p1[v+5][h] - p1[v+3][h]) +
                    (p1[v+5][h+2] - p1[v+3][h+2]) +
20                    2 * (p1[v+5][h+1] - p1[v+3][h+1])) +
                    abs((p1[v+3][h+2] - p1[v+3][h]) +
                    (p1[v+5][h+2] - p1[v+5][h]) +
                    2 * (p1[v+4][h+2] - p1[v+4][h])), 255);
    }
25 }

```

5.1.6—Parameterized Function

Source code

The benchmark source code is not very likely to be written in that form in real world applications. Normally, it would be encapsulated in a function with parameters for input and output arrays along with the sizes of the picture to work on.

Therefore the source code would look similar to the following:

```
void edge3x3(int *p1, int *p2, int HORLEN, int VERLEN)
```

```

{
  for (v=0; v<=VERLEN-3; v++) {
    for (h=0; h<=HORLEN- 3; h++) {
      htmp = (**(p1 + (v+2) * HORLEN + h) - **(p1 + v * HORLEN
5 + h)) +
              (**(p1 + (v+2) * HORLEN + h+2) -**(p1 + v * HORLEN +
h+2)) +
              2 * (**(p1 + (v+2) * HORLEN + h+1) -**(p1 + v * HORLEN
+ h+1));
10      if (htmp < 0)
          htmp = htmp;
          vtmp = (**(p1 + v * HORLEN + h+2) - **(p1 + v * HORLEN +
h)) +
                  (**(p1 + (v+2) * HORLEN + h+2) -**(p1 + (v+2) *
15 HORLEN + h) )+
                  2 * (**(p1 + (v+1) * HORLEN + h+2) -**(p1 + (v+1) *
HORLEN + h));
          if (vtmp < 0)
              vtmp = vtmp;
20      sum = htmp + vtmp;
          if (sum > 255)
              sum = 255;
          ** (p2 + (v+1) * HORLEN + h+1) = sum;
      }
25  }
}

```

This requires some additional features from the compiler.

- ■—interprocedural optimizations and analysis
- 30 • ■—hints by the Programmer, (e.g., a compiler known assert(VERLEN % 2 == 0) makes unroll-and-jam actually possible without peeling off iterations and running them conditionally).

Fitting the Algorithm Optimally to the Array

Since HORLEN and VERLEN are not known at compile time these unknown parameters introduce some constraints which prevent pipeline vectorization. The compiler must assume that the IRAMs cannot hold all HORLEN input values in a row, so pipeline vectorization would not be possible.

Strip Mining Inner Loop

Strip mining partitions the inner loop into a loop that runs over a strip, which is chosen to be of the same size as the IRAMs can hold and a by strip loop iterating over the strips. Of course the strip loops upper bound must be adjusted for the possible incomplete last strip. After the strip mining, the original code would look like be approximately as follows (outer v-loop neglected):

```
for (h=0; h <= HORLEN-3; h+= stripsize)
    for (hh=h; h<=min (h+stripsize--1, HORLEN-3) ; hh++) {
        htmp = ( * * (p1 + (v+2) * HORLEN + hh) - * * (p1 + v *
HORLEN + hh) ) +
        . . .
    }
}
```

Assuming an IRAM size strip size of 256, the following simulated results can be obtained for one strip. The values must be multiplied with the number of strips to be calculated.

Parameter	Value (shift register synthesis)		Value (data duplication – with IRAM placement)	
Vector length	16		16	
Reused data set size	256		256	
I/O IRAMs	4 I + 2 O = 6		12 I + 2 O = 14	
ALU	45		37	
BREG	31 (12 defined + 19 route)		42 (4 defined + 38 route)	
FREG	29 (1 defined + 28 route)		18 (1 defined + 17 route)	
Data flow graph width	14		14	
Data flow graph height	3 (shift registers) + 8 (calculation)		8 (calculation)	
Configuration cycles (simulated)	configuration	2753	configuration	2754
	preloads	7*4*64 1792	preloads	7*12*64 5376
	cycles	128*530 67840	cycles	128*553 70784

	sum	72385	sum	78914
--	-----	-------	-----	-------

The RISC DSP needs about 1.47 million cycles for this amount of data. As mentioned above, these values do not include cache miss penalties and truly underestimate the real values.

Furthermore, it can be seen that data duplication does not improve the performance. The

5 reason for this seems to be ~~ana~~ worse placement and— routing.

5.2—FIR Filter

5.2.1—Original Code

Source code:

```

10  #define N 256
    #define M 8

    for (i = 0; i < N-M+1; i++) {
S:   y[i] = 0;
15     for (j = 0; j < M; j++)
S':  y[i] += c[j] * x[i+M-j-1];
    }

```

The constants N and M are replaced by their values by the pre-processor. The data

20 ~~dependencedependency~~ graph is ~~the following:~~ shown in Fig. 31.

```

for (i = 0; i < 269; i++) {
    S:   y[i] = 0;
        for (j = 0; j < 8; j++)
S':  y[i] += c[j] * x[i+7-j];
25 }

```

~~We have the~~The following is a corresponding table:

Parameter	Value
Vector length	269
Reused data set size	-
I/O IRAMs	3
ALU	2
BREG	0
FREG	0
Data flow graph width	1

Data flow graph height	2
Configuration cycles	2+8=10

5.2.2 First Solution

In ~~the~~ a case we want in which it is desired to save memory, ~~the~~ a straightforward solution is to unroll the inner loop and to use shift register synthesis to delay the values of array x in the pipeline. No other optimization is applied before as either they do not have an effect on the loop or they increase the need for IRAMs. After loop unrolling, ~~we obtain~~ the following code is obtained:

```

for (i = 0; i < 269; i++) {
    Y[i] = 0;
    Y[i] += c[0] * x[i+7];
    Y[i] += c[1] * x[i+6];
    Y[i] += c[2] * x[i+5];
    Y[i] += c[3] * x[i+4];
    Y[i] += c[4] * x[i+3];
    Y[i] += c[5] * x[i+2];
    Y[i] += c[6] * x[i+1];
    Y[i] += c[7] * x[i];
}

```

~~Then the table looks like this:~~

The following is a corresponding table:

Parameter	Value
Vector length	269
Reused data set size	-
I/O IRAMs	9
ALU	16
BREG	0
FREG	0
Data flow graph width	2
Data flow graph height	9
Configuration cycles	9+269=278

Dataflow analysis reveals that $y[0]=f(x[0], \dots, x[7])$, $y[1]=f(x[1], \dots, x[8])$, \dots

$y[i]=f(x[i], \dots, x[i+7])$. Successive values of y depend on almost the same successive

values of x . To prevent unnecessary accesses to the IRAMs, the values of x needed for the

computation of the next values of y are kept in registers. In ~~our~~this case, this shift register synthesis needs 7 registers. This will be achieved on the PACT XPP, by keeping them ~~into~~in FREGs. Then ~~we obtain the dataflow graph depicted below.~~of Fig. 32 is obtained. An IRAM is used for the input values and an IRAM for the output values. The first 8 cycles are used to fill the pipeline and then the throughput is of one output value/cycle. ~~We can depict the~~The code may be represented as the following follows:

```

r0 = x[0];
r1 = x[1];
r2 = x[2];
10  r3 = x[3];
    r4 = x[4];
    r5 = x[5];
    r6 = x[6];
    r7 = x[7];
15  for (i = 0; i < 269; i++) {
        y[i] = c7*r0 + c6*r1 + c5*r2 + c4*r3 + c3*r4 + c2*r5 +
        c1*r6 + c0*r7;
        r0 = r1;
        r1 = r2;
20    r2 = r3;
        r3 = r4;
        r4 = r5;
        r5 = r6;
        r6 = r7;
25    r7 = x[i+7];
    }

```

~~The~~A final table is shown below, and the expected speedup with respect to a standard superscalar processor with 2 instructions issued per cycle is 13.6.

Parameter	Value
Vector length	269
Reused data set size	-
I/O IRAMs	2
ALU	16

BREG	0
FREG	7
Data flow graph width	3
Data flow graph height	9
Configuration cycles	8+269=277

Ops	Number
LD/ST (2 cycles)	2
ADDRCOMP (1 cycle)	0
ADD/SUB (1 cycle)	8
MUL (2 cycles)	-8
SHIFT (1 cycle)	0
Cycles per iteration	28
Cycles needed for the loop (2-way)	(28*269)/2=3766

Variant with Larger Loop Bounds

5 ~~Let us take~~ Taking larger loop bounds and ~~set~~ setting the values of N and M to 1024 and 64.
64:

```

for (i = 0; i < 961; i++)      {
    y[i] = 0;
    for (j = 0; j < 64; j++)
10        y[i] += c[j] * x[i+63-j];
}
```

Following the loop optimizations driver given before, ~~we apply~~ loop tiling is applied to reduce the iteration range of the inner loop. ~~We obtain the~~ The following loop nest is

15 obtained:

```

for (i = 0; i < 961; i++) {
    y[i] = 0;
    for (jj = 0; jj < 8; jj++)
        for (j = 0; j < 8; j++)
20        y[i] += c[8*jj+j] * x[i+63-8*jj-j];
}
```


A subsequent application of loop unrolling on the inner loop yields:

```

for (i = 0; i < 961; i++) {
    y[i] = 0;
    for (jj = 0; jj < 8; jj++) {
5        y[i] += c[8*jj] * x[i+63-8*jj];
        y[i] += c[8*jj+1] * x[i+62-8*jj];
        y[i] += c[8*jj+2] * x[i+61-8*jj];
        y[i] += c[8*jj+3] * x[i+60-8*jj];
        y[i] += c[8*jj+4] * x[i+59-8*jj];
10       y[i] += c[8*jj+5] * x[i+58-8*jj];
        y[i] += c[8*jj+6] * x[i+57-8*jj];
        y[i] += c[8*jj+7] * x[i+56-8*jj];
    }
}

```

- 15 Finally ~~we obtain~~, the same dataflow graph as above is obtained, except that the coefficients must be read from another IRAM rather than being directly handled like constants by the multiplications. After shift register synthesis, the code is may be the following:

```

for (i = 0; i < 961; i++) {
    r0 = x[i+56];
20    r1 = x[i+57];
    r2 = x[i+58];
    r3 = x[i+59];
    r4 = x[i+60];
    r5 = x[i+61];
25    r6 = x[i+62];
    r7 = x[i+63];
    for (jj = 0; jj < 8; jj++)
        y[i] = c[8*jj]*r0 + c[8*jj+1]* r1 + c[8*jj+2]*r2 +
                c[8*jj+3]*r3 + c[8*jj+4]*r4 + c[8*jj+5]*r5 +
30         c[8*jj+6]*r6 + c[8*jj+7]*r7;

    r0 = r1;
    r1 = r2;
    r2 = r3;
    r3 = r4;
}

```

```

r4      = r5;
r5      = r6;
r6      = r7;
r7      = x[i+63-8*jj];
5      }
      }

```

The following table is the same ~~than before~~ as above except for the vector length and the expected speedup with respect to a standard superscalar processor with 2 instructions issued per cycle is 17.5.

Parameter	Value
Vector length	8
Reused data set- size	-
I/O IRAMs	2
ALU	16
BREG	0
FREG	7
Data flow graph width	3
Data flow graph height	9
Configuration cycles	8+8=16

Ops	Number
LD/ST (2 cycles)	10
ADDRCOMP (1 cycle)	0
ADD/SUB (1 cycle)	16
MUL (2 cycles)	17
SHIFT (1 cycle)	0
Cycles per iteration	70
Cycles needed for the loop (2-way)	(70*8)/2=280

15 5.2.3—More Parallel Solution

The solution ~~we presented~~ above does not expose a lot of parallelism in the loop. ~~We can try to~~ To explicitly parallelize the loop before ~~we generate~~ generating the dataflow graph. Of course ~~exposing~~ can be tried. Exposing more parallelism ~~means~~ may mean more pressure on the memory hierarchy.

20

In the data dependence graph presented at the ~~beginning~~ above, the only loop-carried dependence is the dependence on S₂¹ and it is only caused by the reference to y[i]. Hence ~~we~~

apply, node splitting is applied to get a more suitable data dependence graph. We obtain then Accordingly, the following may be obtained:

```

    for (i = 0; i < 249; i++) {
        y[i] = 0;
5       for (j = 0; j < 8; j++)
            {
                tmp = c[j] * x[i+7-j];
                y[i] += tmp;
            }
10    }

```

Then scalar expansion ~~is~~ may be performed on *tmp* to remove the anti loop-carried dependence caused by it, and ~~we have~~ the following code may be obtained:

```

    for (i = 0; i < 249; i++) {
15       y[i] = 0;
        for (j = 0; j < 8; j++)
            {
                tmp[j] = c[j] * x[i+7-j];
                y[i] += tmp[j];
20            }
        }

```

The parameter table is the following:

Parameter	Value
Vector length	249
Reused data set size	-
I/O IRAMs	3
ALU	2
BREG	0
FREG	1
Data flow graph width	2
Data flow graph height	2
Configuration cycles	2+8=10

Then we apply loop distribution may then be applied to get a vectorizable and a not vectorizable loop.

```

    for (i = 0; i < 249; i++) {
        y[i] = 0;
5       for (j = 0; j < 8; j++)
            tmp[j] = c[j] * x[i+7-j];
        for (j = 0; j < 8; j++)
            y[i] += tmp[j];
    } }

```

The following parameter table ~~given below~~ corresponds to the two inner loops in order to be compared with the preceding table.

Parameter	Value
Vector length	249
Reused data set size	-
I/O IRAMs	5
ALU	2
BREG	0
FREG	1
Data flow graph width	1
Data flow graph height	3
Configuration cycles	$1*8+1*8=16$

Then we must take into account the The architecture may be taken into account. The first loop is fully parallel; this, which means that we would need $2*8=16$ input values at a time. This is all right, as it corresponds to the number of IRAMs of the PACT XPP. Hence ~~we do not need~~, to strip-mine the first inner loop. ~~The case of~~ is not required. To strip-mine the second loop is ~~trivial, it does not, need to be strip-mined either, also not required~~. The second loop is a reduction, ~~it~~. It computes the sum of a vector. This ~~is~~ may be easily found by the reduction recognition optimization and ~~we obtain~~ the following code may be obtained.

```

for (i = 0; i < 249; i++) {
    y[i] = 0;
    for (j = 0; j < 8; j++)
25    tmp[j] = c[j] * x[i+7-j];
}

```

```

        /* load the partial sums from memory using a shorter
vector length */
        for (j = 0; j < 4; j++)
            aux[j] = tmp[2*j] + tmp[2*j+1] ;

5
        /* accumulate the short vector */
        for (j = 0; j < 1; j++)=
            aux[2*j] = aux[2*j] + aux[2*j+1];

10
        /* sequence of scalar instructions to add up the partial
sums */
        y[i] = aux[0] + aux(2);
    }

```

- 15 Like above-~~we give~~, only one table is given below for all innermost loops and the last instruction computing $y[i]$.

Parameter	Value
Vector length	249
Reused data set size	-
I/O IRAMs:	12
ALU	4
BREG:	0
FREG	0
Data flow graph width	1
Data flow graph height	4
Configuration cycles	$1*8+1*4+1*1=13$

- 20 Finally, loop- unrolling ~~is~~may be applied on the inner loops,~~the~~. The number of operations is always less than the number of processing elements of the PACT XPP.

```

        for (i = 0; i < 961; i++)
        {
            tmp[0] = c[0] * x[i+7];
            tmp[1] = c[1] * x[i+6];
25
            tmp[2] = c[2] * x[i+5];
            tmp[3] = c[3] * x[i+4];
            tmp[4] = c[4] * x[i+3];

```

```

tmp[5] = c[5] * x[-i+2];
tmp[6] = c[6] * x[i+1];
tmp[7] = c[7] * x[i];

```

```

5      aux[0] = tmp[0] + tmp[1];

      aux[1] = tmp[2] + tmp[3];

      aux[2] = tmp[4] + tmp[5];

      aux[3] = tmp[6] + tmp[7];

10     aux[0] = aux[0] + aux[1];

      aux[2] = aux[2] + aux[3];

      y[i] = aux[0] + aux[2];

      }

```

15

~~We obtain then the following~~The dataflow graph illustrated in Fig. 33, representing the inner loop, may be obtained.

It could be mapped on the PACT XPP with each layer executed in parallel, thus

20 ~~needing~~requiring 4- cycles/iteration and 15 ALU-PAEs, 8 of which are needed in parallel.

As the graph is already synchronized, the throughput reaches one iteration/cycle, after 4 cycles to fill the pipeline. The coefficients are taken —as constant inputs by the ALUs performing the multiplications.

25 ~~The~~A drawback of this solution ~~is~~may be that it uses 16 IRAMs, and that the input data must be stored in a special order. The number of needed IRAMs can be reduced if the coefficients are handled like constant for each ALU. But due to data locality of the program, ~~we~~it can ~~assume~~be assumed that the data already reside in the cache. ~~And as~~As the transfer of data

from the cache to the IRAMs can be achieved efficiently, the configuration can be executed on the PACT XPP without waiting for –the data to be ready in the IRAMs. ~~The~~Accordingly, ~~the~~ parameter table ~~is then~~may be the following:

Parameter	Value
Vector length	249
Reused data set size	-
I/O IRAMs	16
ALU	15
BREG	0
FREG	0
Data flow graph width	8
Data flow graph height	4
Configuration cycles	4+961

5

Variant with Larger Bounds

To make the things a bit more interesting, ~~we set~~in one case, the values of N and M were set to 1024 and 64.

```

10      for (i = 0; i < 961; i++) {
          y[i] = 0;
          for (j = 0; j < 64; j++)
              y[i] += c[j] * x[i+63-j];
      }

```

15 The data dependence graph is the same as above. ~~We apply~~Node splitting may then node-sp
littingbe applied to get a more convenient data dependence graph.

```

20      for (i = 0; i < 961; i++) {
          y[i] = 0;
          for (j = 0; j < 64; j++)
              {
                  tmp = c[j] * x[i+63-j];
                  y[i] += tmp;
              }
      }

```

25

After scalar expansion:

```

    for (i = 0; i < 961; i++) {
        y[i] = 0,
        for (j = 0; j < 64; j++)
5          {
            tmp[j] = c[j] * x[i+63-j];
            y[i] += tmp[j];
          }
    }
10
```

After loop distribution:

```

    for (i = 0; i < 961; i++) {
        y[i] = 0;

15        for (j = 0; j < 64; j++)

            tmp[j] = c[j] * x[i+63-j];

        for (j = 0; j < 64; j++)

            y[i] += tmp[j];

20    }}

```

~~We go~~After going through the compiling process, ~~and we arrive to~~ the set of optimizations that depends upon architectural parameters. ~~We want~~ may be arrived at. It might be desired to split the iteration space, as too many operations would have to be performed in parallel, if ~~we keep it is kept~~ as such. Hence ~~we perform~~, strip-mining may be performed on the 2 loops. ~~We can only access~~ Only 16 data can be accessed at a time, so, because of the first loop, the factor will be $64 * 2/16 = 8$ for the 2 loops (as ~~we always have in mind that we want~~ it is desired to execute both at the same time on the PACT XPP).


```

    for (i = 0; i < 961; i++) {
        y[i] = 0
        for (jj = 0; jj < 8; jj++)
            for (j = 0; j < 8; j++)
5           tmp[8*jj+j] = c[8*jj+j] * x[i+63-8*jj-j];
        for (jj = 0; jj < 8; jj++)
            for (j = 0; j < 8; j++)
                y[i] += tmp[8*jj+j];
    }

```

10

~~Now we apply reduction recognition on the second innermost loop.~~

Then, loop fusion on the *jj* loops may be performed.

```

    for (i = 0; i < 961; i++) {
15       y[i] = 0;
        for (jj = 0; jj < 8; jj++) {
            for (j = 0; j < 8; j++)
                tmp[8*jj+j] = c[8*jj+j] * x[i+63-8*jj-j];
            for (j = 0; j < 8; j++)
20              y[i] += tmp[8*jj+j];
        }
    }

```

Reduction recognition may then be applied on the second innermost loop.

```

25  for (i = 0; i < 961; i++) {
        tmp = 0;
        for (jj = 0; jj < 8; jj++)
        {
            for (j = 0; j < 8; j++)

```

```

    tmp [8*jj+j] = c [8*jj+j] * x[i+63-8*jj-j] ;

    /* load the partial sums from memory using a shorter vector
length */
5      for (j = 0; j < 4; j++)
        aux[j] = tmp [8*jj+2*j] + tmp [8*jj+2*j+1];

    /* accumulate the short vector */
        for; (j = 0; j < 1; j++)
10     aux[2*j] = aux [2*j] + aux[2*j+1];

    /* sequence of scalar instructions to add up the partial
sums */
        y[i] = aux[0] + aux[2] ;

15
And then loopLoop unrolling may then be performed:
    for (i = 0; i < 961; i++)
        for (jj = 0; jj < 8; jj++)
        {
20     tmp[8*jj] = c[8*jj] * x[i+63-8*jj];
            tmp[8*jj+1] = c[8*jj+1] * x[i+62-8*jj];
            tmp[8*jj+2] = c[8*jj+2] * x[i+61-8*jj];
            tmp[8*jj+3] = c[8*jj+3] * x[i+59-8*jj];
            tmp[8*jj+4] = c[8*jj+4] * x[i+58-8*jj];
25     tmp[8*jj+5] = c[8*jj+5] * x[i+57-8*jj];
            tmp[8*jj+6] = c[8*jj+6] * x[i+56-8*jj];
            tmp[8*jj+7] = c[8*jj+7] * x[i+55-8*jj];

            aux [0] = tmp[8*jj] 4 + tmp[8*jj+1];
30     aux [1] = tmp [8*jj+2]      + tmp [8*jj+3] ;
            aux [2] = tmp[8*jj+4]      + tmp[8*jj+5];
            aux [3] = tmp [8*jj+6]      + tmp [8*jj+7];

            aux [0] = aux [0] + aux [1];

```

```
aux[2] = aux[2] + aux[3];
```

```
y[i] = aux[0] + aux[2];
```

```
}
```

5

~~We implement the~~The innermost loop may be implemented on the PACT XPP directly with a counter. The IRAMs ~~are~~may be used in FIFO mode, and filled according to the addresses of the arrays in the loop. IRAM0, IRAM2, IRAM4, IRAM6 and IRAM8 contain array 'c'.

10 IRAM1, IRAM3, IRAM5 and IRAM7 contain array 'x'. Array 'c' contains 64 elements, ~~that~~ is i.e., each IRAM contains 8 elements. Array 'x' contains 1024 elements, ~~that is i.e.~~, 128 elements for each IRAM. Array 'y' is directly written to memory, as it is a global array and its address is constant. This constant is used to initialize the address counter of the configuration. ~~The~~A final parameter table is the following:

Parameter	Value
Vector length	8
Reused data set size	-
I/O IRAMs	16
ALU	15
BREG	0
FREG	0
Data flow graph width	8
Data flow graph height	4
Configuration cycles	4+8=12

15

Nevertheless, it should be noted that this version should be less efficient than the previous one. As the same data must be loaded in the different IRAMs from the cache, ~~we have there~~ are a lot of transfers to ~~achieve~~be achieved before the configuration can begin the computations. This overhead must be taken into account by the compiler when choosing the code generation strategy. This means also that the first solution is the solution that will be chosen by the compiler.

20

5.2.4—Other Variant

Source Code

```
25 for (i = 0; i < N-M+1; i++) {
    tmp = 0;
    for (j = 0; j < M; j++)
```

```

        tmp += c[j] * x[i+M-j-1];
        x[i] = tmp;
    }

```

- 5 In this case, ~~it is trivial that~~ the data dependence graph is cyclic due to dependences on *tmp*. Therefore, scalar expansion is applied on the loop, and ~~we obtain~~, in fact, the ~~same~~ code as the first version of the FIR filter is obtained as shown below.

```

        for (i = 0; i < N-M+1; i++) {
            tmp[i] = 0;
10         for (j = 0; j < M; j++)
            tmp[i] += c[j] * x[i+M-j-1];
            x[i] = tmp[i];
        }

```

15 **5.3—Matrix Multiplication**

5.3.1—Original Code

Source code:

```

#define L 10
#define M 15
20 #define N 20
int A[L][M];
int B[M][N];
int R[L][N];

25 main() {
    int i, j, k, tmp, aux;

    /* input A (L*M values) */
    for (i=0; i<L; i++)
30     for (j=0; j<M; j++)
        scanf("%d", &A[i][j]);

    /* input B (M*N values) */
    for (i=0; i<M; i++)

```

```

        for (j=0; j<N; j++) {
            scanf("%d", &B[i][j]);

/* multiply */
5   for (i=0; i<L; i++)
        for (j=0; j<N; j++) {
            aux = 0;
            for (k=0; k<M; k++)
                aux += A[i][k] * B[k][j];
10          R[i][j] = aux;
        }

/* write data to stream */
        for (i=0; i<L; i++)
15      for (j=0; j<N; j++)
            printf("%d\n", R[i][j]);
    }

```

5.3.2 Preliminary Transformations

20 Since no inline-able function calls are present, no interprocedural code movement is done.

Of the four loop nests, the one with the “/* multiply */” comment is the only candidate for running partly on the XPP. All others have function calls in the loop body and are therefore discarded as candidates very early in the compiler.

25

Dependency Analysis

```

        for (i=0; i<L; i++)
            for (j=0; j<N; j++) {
S1          aux = 0;
30          for (k=0; k<M; k++)
S2              aux += A[i][k] * B[k][j];
S3          R[i][j] = aux;
            }

```

~~Figure 60 Data~~ Fig. 34 shows a data dependency graph for matrix multiplication. The data dependency graph shows no dependencies that prevent pipeline vectorization. The loop carried true dependence from S2 to itself can be handled by a feedback of aux as described in [1].—Markus Weinhardt et al., “Memory Access Optimization for Reconfigurable Systems,” *supra*.

Reverse Loop-Invariant Code Motion

To get a perfect loop nest ~~we move~~, S1 and S3 may be moved inside the k-loop. Therefore, appropriate guards ~~are~~may be generated to protect the assignments. The code after this transformation ~~looks like~~is as follows:

```
for (i=0; i<L; i++)
    for(j=0; j<N; j++)
        for (k=0; k<M; k++) {
            if (k == 0) aux[j] = 0;
            aux[j] += A[i] [k] * B [k] [j];
            if (k == M-1) R[i] [j] = aux[j];
        }
```

Scalar Expansion

Our ~~A~~ goal ~~is~~may be to interchange the loop nests to improve the array accesses to utilize the cache best. ~~Unfortunately~~However, the guarded statements involving ‘aux’ may cause backward loop carried anti-~~dependences~~dependencies carried by the j loop. Scalar expansion ~~will~~may break these ~~dependences~~dependencies, allowing loop interchange.

```
for (i=0; i<L; i++ )
    for (j=0 ; j<N;      j++)
        for (k=0; k<M; k++) {
            if(k == 0) aux[j] = 0;

            aux[j] += A[i] [k] * B[k] [j];

            if (k == M-1) R[i] [j] = aux [j];
        }
```

Loop Interchange for Cache Reuse

Visualizing the main loop shows the iteration spaces for the array accesses. ~~(Figure 61)~~ Fig. 35 is a visualization of array access sequences. Since C arrays are placed in row major order, the cache lines are placed in the array rows. At first sight, there seems to be no need for optimization because the algorithm requires at least one array access to ~~striderstride~~ stride over a column. Nevertheless, this assumption misses the fact that the access rate is of interest, too. Closer examination shows that array R is accessed in every j iteration, while B is accessed every k- iteration, always producing a cache miss². ~~(“aux” is not currently discussed since it is not expected that it would be written to or read from memory, as there are no defs or uses outside the loop nest.)~~ This leaves a possibility for loop interchange to improve cache access as proposed by Kennedy and Allen in [7]. ~~Markus Weinhardt et al., “Pipeline Vectorization,” supra.~~

Figure 61 The visualized array access sequences.

- 15 ~~Finding~~To find the best loop nest is relatively simple. ~~The, the~~ algorithm simply interchanges ~~may interchange~~ each loop of the nests into the innermost position and ~~annotates~~annotate it with the so-called innermost- memory cost term. This cost term is a constant for known loop bounds or a function of the loop bound for unknown loop bounds. The term ~~is~~may be calculated in three steps.
- 20 • ~~■~~—First, the cost of each reference³ in the innermost loop body ~~is~~may be calculated to:
- ~~■~~ 1, if the reference does not depend on the loop induction variable of the (current) innermost loop;
 - ~~■~~ the loop count, if the reference depends on the loop induction variable and strides over a non-contiguous area ~~in~~with respect of the cache layout;
- 25 • ~~■~~ $\frac{N \cdot s}{b}$, if the reference depends on the loop induction variable and strides over a contiguous dimension. In this case, N is the loop count, s is the step size and b is the cache line size, respectively.

² ~~_____ We neglect “aux” in this observation since we do not expect it to be written to or read from memory (no defs or uses outside the loop nest)~~

³ ~~_____ Reference means access to an array in this case. Since the transformation wants to optimize cache access, it must address references to the same array within small distances as one. This prohibits over-estimation of the actual costs.~~

In this case, a “reference” is an access to an array. Since the transformation attempts to optimize cache access, it must address references to the same array within small distances as one. This may prohibit over-estimation of the actual costs.

- 5 • ■—Second, each reference cost ismay be weighted with a factor for each other loop, which is:

- ■ 1, if the reference does not depend on the loop index;
- ■ the loop count, if the reference depends on the loop index.

- 10 • ■—Third, the overall loop nest cost ismay be calculated by summing the costs of all reference costs:

After invoking this algorithm for each loop as the innermost, the one with the lowest cost ismay be chosen as the innermost, the next as the next outermost, and so on.

Innermost loop	R[i][j]	A[i][k]	B[k][j]	Memory access cost
k	$1 \cdot L \cdot N$	$\frac{M}{b} \cdot L$	$M \cdot N$	$L \cdot N + \frac{M}{b} \cdot L + M \cdot N$
i	$1 \cdot L \cdot N$	$1 \cdot L \cdot M$	$1 \cdot M \cdot N$	$L \cdot N + L \cdot M + M \cdot N$
j	$\frac{N}{b} \cdot L$	$L \cdot M$	$\frac{N}{b} \cdot M$	$\frac{N}{b} (L + M) + L \cdot M$

15

Table 1 Loop memory access costs for the different loops being innermost

The preceding table shows the values for the matrix multiplication. Since the j term is the smallest (~~of course~~ assuming $b > 1$), the j-loop is chosen to be the innermost. The next outer loop then is k, and the outermost is i. Thus, the resulting code after loop interchange ismay

20 be:

```

for (i=0; i<L; i++)
    for (k=0; k<M; k++) ,
        for (j=0; j<N; j++) {
            if (k == 0) aux[j] = 0;
            aux[j] += A[i][k] * B[k][j];
            if (k == M-1) R[i][j] = aux[j];
        }

```

25

~~Figure 62 The visualized~~ Fig. 36 shows the improved iteration spaces. It shows array access sequences after optimization. ~~Here the~~ The improvement is visible to the naked eye; since array B is now read following the cache lines. ~~Figure 62 shows the improved iteration~~ spaces. It is to say that this This optimization does not optimize primarily for the XPP; but mainly optimizes the cache-hit rate, thus improving the overall performance.

Unroll and Jam

After improving the cache access behavior, the possibility for reduction recognition has been destroyed. This is a typical example for transformations where one excludes the other. Nevertheless ~~we obtain~~, more parallelism may be obtained by doing unroll-and-jam. Therefore ~~we unroll~~, the outer loop may be partially unrolled with the unroll factor. This factor is mainly chosen by the minimum of two calculations:

- # available IRAMs / # used IRAMs in the inner loop body
- # available ALU resources / # used ALU resources in the inner loop.

In this example embodiment, the accesses to “A” and “B” depend on k (the loop which will be unrolled). Therefore, they must be considered in the calculation. The accesses to “aux” and “R” do not depend on k. Thus, they can be subtracted from the available IRAMs, but ~~do not~~ need to be added to the denominator. Therefore ~~we calculate~~, (assuming an XPP64) $14/2 = 7$ is calculated for the unroll factor obtained by the IRAM resources.

On the other hand, the loop body involves two ALU operations (1 add, 1 mult), which may yield an unrolling factor of approximately $64/2 = 32$ ^[4].

⁴ ~~This is a very inaccurate estimation, since it neither estimates the resources spent by the controlling network, which decreases the unroll factor, nor takes it into account that e.g. the BREG-PAEs also have an adder, which increases the unroll factor. Although it has no influence to this example the unroll factor calculation of course has to account for this in a production compiler.~~

(This is an inaccurate estimation since it neither estimates the resources spent by the controlling network, which may decrease the unroll factor, nor takes into account that, e.g., the BREG-PAEs also have an adder, which may increase the unroll factor. Although it does not influence this example, the unroll factor calculation should account for this in a production compiler.) The constraint generated by the IRAMs therefore dominates by far.

Having chosen the unroll factor ~~we must trim our~~, the loop trip count is trimmed to be a multiple of that factor. Since the k loop has a loop count of 15, ~~we peel off~~ the first iteration may be peeled off and unroll—the remaining loop may be unrolled.

```

for (i=0; i<L; i++) {
5      for (k=0; k<1; k++) {
          for (j=0; j<N; j++) {
              if (k==0) aux[j] = 0;
              aux[j] += A[i][k] * B[k][j];
              if (k==M-1) R[i][j] = aux[j];
10          }
      }

      for (k=1; k<M; k+=7) {
          for (j=0; j<N; j++) {
              if (k==0) aux[j] = 0;
15          aux[j] += A[i][k] * B[k][j];
              if (k==M-1) R[i][j] = aux[j];
          }

          for (j=0; j<N; j++) {
              if (k+1==0) aux[j] = 0;
20          aux[j] += A[i][k+1] * B[k+1][j];
              if (k+1==M-1) R[i][j] = aux[j];
          }

          for (j=0; j<N; j++) {
              if (k+2==0) aux[j] = 0;
25          aux[j] += A[i][k+2] * B[k+2][j];
              if (k+2==M-1) R[i][j] = aux[j];
          }
      }
  }

```

```

    }

    for (j=0; j<N; j++) {

        if (k+3==0) aux[j] = 0;

        aux[j] += A[i][k+3] * B[k+3][j];

5         if (k+3==M-1) R[i][j] = aux[j];

    }

    for (j=0; j<N; j++) {

        if (k+4==0) aux[j] = 0;

        aux[j] += A[i][k+4] * B[k+4][j];

10        if (k+4==M-1) R[i][j] = aux[j];

    }

    for (j=0; j<N; j++) {

        if (k+5==0) aux[j] = 0;

        aux[j] += A[i][k+5] * B[k+5][j];

15        if (k+5==M-1) R[i][j] = aux[j];

    }

    for(j=0; j<N; j++) {

        if (k+6==0) aux[j] = 0;

        aux[j] += A[i][k+6] * B[k+6][j];

20        if(k+6==M-1) R[i][j] = aux[j];

    }

}

}

```

Due to ~~the fact that~~ placement by the reverse loop invariant code motion ~~placed~~ of the loop invariant code into the inner loop, which is ~~now~~ duplicated seven times, it is very likely that dead code elimination can get rid of some of these duplicates. Thus, the code ~~is~~ may be shortened to:

```

5  for (i=0; i<L; i++) {
    for (k=0; k<1; k++) {
      for(j=0; j<N; j++) {
        if (k==0) aux[j] = 0;
        aux[j] += A[i][k] * B[k][j];
10      }
    }
    for (k=1; k<M; k+=7) {
      for (j=0; j<N; j++) {
        aux[j] += A[i][k] * B[k][j];
15      }
      for (j=0; j<N; j++) {
        aux[j] += A[i][k+1] * B[k+1][j];
      }
      for(j=0; j<N; j++) {
20        aux[j] += A[i][k+2] * B[k+2][j];
      }
      for (j=0; j<N; j++) {
        aux[j] += A[i][k+3] * B[k+3][j];
      }
      for (j=0; j<N; j++) {
25        aux[j] += A[i][k+4] * B[k+4][j];
      }
      for (j=0; j<N; j++) {
        aux[j] += A[i][k+5] * B[k+5][j];
30      }
      for (j=0; j<N; j++) {
        aux[j] += A[i][k+6] * B[k+6][j];
        if (k+6==M-1) R[i][j]

```

```

    }
}
}

```

5 Before ~~we jam~~jamming of the inner loops we have to, it may be taken into account for the
~~fact that the first iteration of the k loop was peeled of which would produce an own~~
configuration. Since ~~we calculated the unroll-and-jam factor~~is calculated to fit into one
configuration, this side effect ~~has to~~should be prevented. Because it should be no problem to
run the k loop with— variable step sizes, ~~we fuse the- k loops~~may be fused again ~~and adjust,~~
10 the step size ~~and guard~~may be adjusted, and the statements may be guarded. This ~~yields~~may
yield:

```

for (i=0; i<L; i++) {
    for (k=0; k<M; k+= k<1 ? 1 : 7) {
        for (j=0; j<N; j++) {
15             if (k==0) aux[j] = 0;
                if (k==0) aux[j] += A[i][k] * B[k][j];
        }
        for(j=0; j<N; j++) {
            if (k>0) aux[j] += A[i][k] * B[k][j];
20        }
        for(j=0; j<N; j++) {
            if (k>0) aux[j] += A[i][k+1] * B[k+1][j];
        }
        for(j=0; j<N; j++) {
25             if (k>0) aux[j] += A[i][k+2] * B[k+2][j];
        }
        for(j=0; j<N; j++) {
            if (k>0) aux[j] += A[i][k+3] * B[k+3][j];
        }
        for(j=0; j<N; j++) {
30             if (k>0) aux[j] += A[i][k+4] * B[k+4][j];
        }
        for(j=0; j<N; j++) {
            if (k>0) aux[j] += A[i][k+5] * B[k+5][j];
        }
    }
}

```

```

    }
    for(j=0; j<N; j++) {
        if (k>0) aux[j] += A[i][k+6] * B[k+6][j];
        if (k+6==M-1) R[i][j] = aux[j];
5      }
    }
}

```

Now ~~we can jam~~, the inner loops ~~and finally obtain~~ may be jammed, and the following may be obtained.

```

10  for (i=0; i<L; i++) {
    for (k=0; k<M; k+= k<1 ? 1 : 7) {
        for (j=0; j<N; j++) {
            if (k==0) aux[j] = 0;
            if (k==0) aux[j] += A[i][k] * B[k][j];
15          if (k>0) {
                aux[j] += A[i][k] * B[k][j];
                aux[j] += A[i][k+1] * B[k+1][j];
                aux[j] += A[i][k+2] * B[k+2][j];
20          aux[j] += A[i][k+3] * B[k+3][j];
                aux[j] += A[i][k+4] * B[k+4][j];
                aux[j] += A[i][k+5] * B[k+5][j];
                aux[j] += A[i][k+6] * B[k+6][j];
                if (k+6==M-1) R[i][j] = aux[j];
25          }
        }
    }
}

```

30 5.3.3—XPP Code Generation

The innermost loop can be synthesized in a configuration, which uses 14 IRAMs for the input data, one IRAM to temporary store aux, and one IRAM for the output may R. Furthermore, it ~~is~~ may be necessary to pass the value of k to the XPP to direct the dataflow. This may be done by a streaming input. ~~Figure 63~~ Fig. 37 shows the dataflow graph of the synthesized

configuration. ~~Figure 63 Dataflow graph of~~ and shows matrix multiplication after unroll and jam. The rightmost 3 branches are omitted. ~~Event from the graph and event connections are emphasized in red color highlighted.~~

5 The following code shows the pseudo code that may be executed on the RISC processor.

```

XPPPreload( config )
for (i=0; i<L; i++) {
    XPPPreload(0, &A[i][0], M)

    XPPPreload(1, &A[i][0], M)

10  XPPPreload(2, &A[i][0], M)

    XPPPreload(3, &A[i][0], M)

    XPPPreload(4, &A[i][0], M)

    XPPPreload(5, &A[i][0], M)

    XPPPreload(6, &A[i][0], M)

15  XPPPreloadClean(15, &R[i][0], M)

    for (k=0; k<M; k+= k<1 ? 1 : 7) {
        XPPPreload(7, &B[k][0], N)
        XPPPreload(8, &B[k+1][0], N)
        XPPPreload(9, &B[k+2][0], N)
20  XPPPreload(10, &B[k+3][0], N)
        XPPPreload(11, &B[k+4][0], N)
        XPPPreload(12, &B[k+5][0], N)
        XPPPreload(13, &B[k+6][0], N)
        XPPEecute(config, IRAM(0), IRAM(1), IRAM(2),
25  IRAM(3), IRAM(4), IRAM(5),
        IRAM(6), IRAM(7), IRAM(8),
        IRAM(9), IRAM(10), IRAM(11),
        IRAM(12), IRAM(13), IRAM(15), k)

    }

30 }

```

The following table shows the simulated configuration. The complete multiplication needs about 3120 cycles without the preloading and configuration. A typical RISC-DSP core with two MAC units and ~~hard-ware~~hardware loop support needs over 26000 cycles (when data is in zero-latency internal memory). Although the time for preloads and cache misses is neglected here, the values ~~promise~~according to an embodiment of the present invention may result in improvements of 200-300 percent compared to a standalone RISC core.

The following is a corresponding parameter table.

Parameter	Value		
Vector length	20		
Reused data set size	20		
I/O IRAMs	14 I + 1 O + 1 internal		
ALU	20		
BREG	26 (8 defined + 18 route)		
FREG	28 (4 defined + 24 route)		
Data flow graph width	14		
Data flow graph height	6 (without routing and balancing)		
Configuration cycles (simulated)	configuration		2633
	preloads	$10 \times 3 \times 7 \times 5$	1050
	cycles	$10 \times 7 \times 15$	1050
	sum	$(k=0) 112 +$ $(k=1) 100 +$ $(k=7) 100$ $\times 10 =$ 3120	7853

5.4—Viterbi Encoder

5.4.1—Original Code

Source Code:

```
/* C-language butterfly */
```

```
#define BFLY(i) {\
```

```
    unsigned char metric, m0, m1, decision; \
```

```
        metric = ((Branchtab29_1[i] ^ sym1) +  
                  (Branchtab29_2 [i] ^ sym2) + 1)/2; \
```

```
        m0 = vp->old_metrics [i] + metric; \
```

```
        m1 = vp->old_metrics [i+128] + (15 - metric) ; \
```



```

    decision = (m0-m1) >= 0; \
    vp->new_metrics [2*i] = decision ? m1 : m0; \
    vp->dp->w[i/16] |= decision << ((2*i)      &31) ; \
    m0 -= (metric+metric-15); \
5    m1 += (metric+metric-15); \

    decision = (m0-m1) >= 0; \

    vp->new_metrics [2*i+1] = decision ? m1 : m0; \

    vp->dp->w[i/16] |= decision << ((2*i+1)&31); \

}

10

int update_viterbi29(void *p,unsigned char sym1,unsigned char
sym2) {

    int i;

    struct v29 *vp = p;

15    unsigned char *tmp;

    int normalize = 0;

    for (i= 0; i<8 ; i++)

        vp->dp->w [i] = 0 ;

20

    for (i =0; i<128; i++)

        BFLY(i);

    /* Renormalize metrics */

25    if (vp->new_metrics[0] > 150) {

```

```

    int i;

    unsigned char minmetric = 255;

    for (i=0; i<64; i++)
5         if (vp->new _metrics [i] < minmetric)

            minmetric = vp->new _metrics [i];

    for (i=0; i<64; i++)

        vp->new _metrics [i]    -= minmetric;

    normalize = minmetric;
10    }

    vp->dp++;

    tmp = vp->old _metrics;

    vp->old_metrics = vp->new _metrics;
15    vp->new_metrics = tmp;

    return normalize;

}

```

20 **5.4.2—Interprocedural Optimizations and Scalar Transformations**

Since no inline-able function calls are present, in an embodiment of the present invention, no interprocedural code movement is done.

After expression simplification, strength reduction, SSA renaming, copy coalescing and
 25 idiom recognition, the code ~~looks like~~ may be approximately as presented below (statements

are reordered for convenience). Note that idiom recognition ~~will~~may find the combination of *min()* and use of the comparison result for *decision* and *_decision*. However, the resulting computation cannot be expressed in C, so ~~we describe it~~ is described below as a comment:

```

int update_viterbi29 (void *p,unsigned char sym1,unsigned char
5 sym2) {
    int i;
    struct v29 *vp = p;
    unsigned char *tmp;
    int normalize = 0;
10
    char *_vpdpw = vp->dp->w;
    for (i=0 ; i<8; i++ )
        * _vpdpw _++ = 0;
    char *_bt29_1= Branchtab29_1;
15 char *_bt29_2= Branchtab29_2;
    char *_vpom0= vp->old_metrics;
    char *_vpom128= vp->old_metrics+128;
    char * vpm= vp->new_metrics;
    char *_vpdpw= vp->dp->w;
20
    for (i=0; i<128; i++) {
        unsigned char metric, _tmp, m0, m1, _m0, _m1, decision,
        _decision;
25
        metric = ((* _bt29_1++ ^ sym1) +
                (* _bt29_2++ ^ sym2) + 1)/2;
        _tmp= (metric+metric-15);
        m0 = *_vpom++ + metric;

```

```

    m1 = *_vpom128++ + (15 - metric);

    _m0 = m0 - _tmp;

    _m1 = m1 + _tmp;

    // decision = m0 >= m1;

5    // _decision = _m0 >= m1;

    *_vpnm++ = min (m0, m1);          // = decision ? m1 : m0

    *_vpnm++ = min(_m0, _m1);        // = _decision ? _m1 : _m0

    _vpdpw[i >> 4] |= ( m0 >= m1) /* decision*/ << ((2*i) &
31)

10    | (_m0 >= _m1) /*_decision*/ <<
    ((2*i+1)&31);
    }

    /* Renormalize metrics */
15    if (vp->new _metrics [0] > 150) {
        int i;
        unsigned char minmetric = 255;

        char *_vpnm= vp->new _metrics;
20        for (i=0; i<64; i++)
            minmetric = min (minmetric, *_vpnm++);

        char *_vpnm= vp->new _metrics;
        for (i=0; i<64; i++)
25            *_vpnm++ -= minmetric;
        normalize = minmetric;
    }

    vp->dp++;

```

```

    tmp = vp->old _metrics;

    vp->old _metrics = vp->new _metrics;

    vp->new_metrics = tmp;

5    return normalize;

}

```

5.4.3—Initialization

The first loop (setting vp->dp->w[0..7] to zero) ~~is~~may be most efficiently executed on
10 the RISC.

5.4.4—Butterfly Loop

The second loop (with the *BFLY()* macro expanded) is of interest for the XPP compiler and needs further examination:

```

15    char                *iram0=    Branchtab29_1;        //
XPPPreload (0, Branchtab29 _1, 128/4) ;
    char    *iram2=    Branchtab29_2;                // XPPPreload (2,
Branchtab29_2, 128/4) ;
    char    *iram4=    vp->old_metrics; // XPPPreload (4, vp-
20 >old_metrics,    128/4);
    char    *iram5=    vp->old _metrics+128; // XPPPreload (5,
vp->old _metrics+128,128/4) ;
    short    *iram6=    vp->new_metrics; // XPPPreload (6, vp-
>new _metrics, 128/2);
25    unsigned long *iram7= vp->dp->w; // XPPPreload (7, vp->dp-
>w, 8);
// sym1 & sym2 are in IRAM 1 & 3

    for (i=0; i<128; i++) {
30        unsigned char metric, _tmp, m0, m1, _m0, _m1

```

```

    metric = ((*iram0++ ^ sym1) +
              (*iram1++ ^ sym2) + 1) / 2;
    _tmp= (metric << 1) -15;
    m0 = *iram2++ + metric;
5    m1 = *iram3++ + (15 - metric) ;
    _m0 = m0 - _tmp;
    _m1 = m1 + _tmp;
    // assuming big endian; little endian has the shift on
the latter min()
10    *iram6++ = (min(m0, m1) << 8) | min (_m0, _m1) ;
    *iram7 [i >> 4] |= (m0 >= m1) << ( (2*i) &
31)
    | (_m0 >= _m1) << ((2*i+1) &31) ;
    }
15

```

The corresponding data flow graph is as followsshown in Fig. 38 (for now ignoring the fact, that the IRAM accesses are mostly char accesses). The solid lines represent data flow, while the dashed lines represent event flow.

20 The following is a corresponding parameter table.

Parameter	Value
Vector length	128
Reused data set size	-
I/O IRAMs	6 I + 2 O
ALU	25
BREG	few
FREG	few
Data flow graph width	4
Data flow graph height	11
Configuration cycles	11+128

25 ~~We~~Some problems are immediately see some problemsnoticed: IRAM7 is fully busy reading and rewriting the same address sixteen times. Loop tiling to a tile size of sixteen gives the *redundant load store elimination* a chance to read the value once and accumulate the bits in ~~the temporary~~temporarily, writing the value to the IRAM at the end of this inner loop. Loop Fusion with the initialization loop then ~~allows~~may allow propagation of the zero values set in

the first loop to the reads of `vp->dp->w[i]` (IRAM7), eliminating the first loop altogether. Loop tiling with a tile size of 16 may also eliminate the `& 31` expressions for the shift values. Since the new inner loop only runs from 0 to 16, the value range analysis now finds that the `& 31` expression is not limiting the value range any further.

5

All remaining input IRAMs are character (8 bit) based. So ~~we need it~~ it may be required for split networks to split the 32-bit stream into four 8-bit streams which are then merged. This adds 3 shifts, 3 ands, and 3 merges for every character IRAM. The merges could be eliminated; when unrolling the loop body. However, unrolling ~~is~~ may be limited to unrolling twice due to ALU availability as well as due to ~~the fact,~~ that IRAM6 is already 16 bit based: ~~unrolling.~~ Unrolling once requires a *shift by 16* and an *or* to write 32 bits in every cycle; ~~unrolling.~~ Unrolling further cannot increase pipeline throughput any more. So the body is only unrolled once, eliminating one layer of merges. This ~~yields~~ may yield two separate pipelines; that each handle two eight bit slices of the 32-bit value from the IRAM, serialized by merges.

10

15

The modified code ~~now looks like~~ may be approximately as follows (unrolling and splitting omitted for simplicity):

```

char                                *iram0= Branchtab29_1;    //
20  XPPPreload(0, Branchtab29_1, 128/4);
    char    *iram2= Branchtab29_2;                // XPPPreload(2,
Branchtab29_2, 128/4);
    char    *iram4= vp->old_metrics; // XPPPreload(4, vp->old
metrics, 128/4);
25  char    *iram5= vp->old_metrics+128; // XPPPreload(5, vp-
>old_metrics+128,128/4);
    short *iram6= vp->new_metrics; // XPPPreload(6, vp-
>new_metrics, 128/2);
    unsigned long *iram7= vp->dp->w; // XPPPreload(7, vp->dp->w,
30  8);
    // sym1 & sym2 are in IRAM 1 & 3

    for (_i=0;_i<8;_i++) {
```

```

    rlse= 0;
    for (i2=0; i2<32; i2+=2) {
        unsigned char metric, _tmp, m0, m1, _m0, _m1;

5         metric = ((*iram0++ ^ sym1) +
                    (*iram1++ ^ sym2) + 1)/2;
        _tmp= (metric << 1) -15;
        m0 = *iram2++ + metric;
        m1 = *iram3++ + (15 - metric);
10        _m0 = m0 - _tmp;
        _m1 = m1 + _tmp;
        *iram6++ = (min(m0,m1) << 8) | min(_m0,_m1);
        rlse = rlse | ( m0 >= m1) << i2
                    | (_m0 >= _m1) << (i2+1);
15    }
    *iram7++ = rlse;
}

```

20 The modified data flow graph (unrolling and splitting omitted for simplicity): ~~And here the is~~
shown in Fig. 39. The splitting network for one IRAM: the is shown in Fig. 40. The bottom
 most level merge is omitted for each level of unrolling.

The following is a corresponding parameter table.

Parameter	Value
Vector length	128
Reused data set size	-
I/O IRAMs	6 I + 2 O
ALU	$2*24+4*3(\text{split})+2(\text{join}) = 62$
BREG	few
FREG	few
Data flow graph width	4
Data flow graph height	11+3 (split)
Configuration cycles	14+64

25 **5.4.5—Re-Normalization:**

The Normalization consists of a loop scanning the input for the minimum and a second loop
 that subtracts the minimum from all elements. There is a data dependency between all

iterations of the first loop and all iterations of the second loop. Therefore, the two loops cannot be merged or pipelined. They ~~will~~may be handled individually.

Minimum Search

- 5 The third loop is a minimum search on a byte array.

```
char *iram0 = vp->new_metrics; // XPPPzeload(0, vp->new_metrics, 64/4);
for (i=0; i<64 ; i++)
    minmetric = min (minmetric, *iram0++);
```

10

The following is a corresponding parameter table.

Parameter	Value
Vector length	64
Reused data set size	-
I/O IRAMs	1+1
ALU	1
BREG	0
FREG	0
Data flow graph width	1
Data flow graph height	1
Configuration cycles	64

Reduction recognition ~~eliminates~~may eliminate the dependence for *minmetric*, enabling a four-times unroll to utilize the IRAM width of 32 bits. A split network has to be added to separate the 8 bit streams using 3 SHIFT and 3 AND operations. Tree balancing may re-
 15 distributesdistribute the *min()* operations to minimize the tree height.

```
char *iram0 = vp->new_metrics; // XPPPpreload (0, vp->new_metrics, 16);
for (i=0; i<16; i++)
  20    minmetric = min(minmetric, min(min(*iram0++, *iram0++),
                                     min(*iram0++, *iram0++)));
```

The following is a corresponding parameter table.

Parameter	Value
Vector length	16
Reused data set size	-
I/O IRAMs	1 I + 1 O

ALU	4*min
BREG	3*shln+3*shrn
FREG	0
Data flow graph width	4
Data flow graph height	5
Configuration cycles	5+16

Reduction recognition again ~~eliminates~~may eliminate the loop carried dependence for *minmetric*, enabling loop tiling and then unroll and jam to increase parallelism; ~~the~~. The maximum for the tiling size is 16 IRAMs / 2 IRAMS = 8. Constant propagation and tree

5 rebalancing ~~reduces~~may reduce the dependence height of the final merging expression:

```

char                                *iram0= vp->new _metrics;
// XPPPreload(0, vp->new _metrics, 2);
char *iram1= vp->new _metrics+8; // XPPPreload(1, vp-
>new _metrics+8, 2);
10 char *iram2= vp->new _metrics+16; // XPPPreload(2, vp-
>new _metrics+16, 2);
char *iram3= vp->new _metrics+24; // XPPPreload(3, vp-
>new _metrics+24, 2);
char                                *iram4= vp->new _metrics+32;
15                                /
/ XPPPreload(4, vp->new _metrics+32, 2);

char *iram5= vp->new _metrics+40; // XPPPreload ( 5, vp-
>new _metrics+40, 2) ;

char *iram6= vp->new _metrics+48; // XPPPreload (6, vp->new
20 _metrics+48, 2) ;

char *iram7= vp->new _metrics+56; // XPPPreload(7, vp-
>new _metrics+56, 2) ;

for (i =0; _i<2; i++) {
25     minmetric0 = min (minmetric0, min( min(*iram0++,
*iram0++),
min(*iram0++, *iram0++)));

```

```

        minmetric1 = min (minmetric1, min( min(*iram1++,
*iram1++),
                                min(*iram1++, *iram1++)));
        minmetric2 = min (minmetric2, min( min(*iram2++,
5  *iram2++),
                                min(*iram2++, *iram2++)));
        minmetric3 = min (minmetric3, min( min(*iram3++,
*iram3++),
                                min(*iram3++, *iram3++)));
10    minmetric4 = min (minmetric4, min( min(*iram4++,
*iram4++),
                                min(*iram4++, *iram4++)));
        minmetric5 = min (minmetric5, min( min(*iram5++,
*iram5++),
15    min(*iram5++, *iram5++)));
        minmetric6 = min (minmetric6, min( min(*iram6++,
*iram6++),
                                min(*iram6++, *iram6++)));
        minmetric7 = min (minmetric7, min( min(*iram7++,
20  *iram7++),
                                min(*iram7++, *iram7++)));
    }
    minmetric = min(min((    min(minmetric_0, minmetric_1),
                            min(minmetric_2, minmetric_3)),
25    min(( min(minmetric_4, minmetric_5),
            min(minmetric_6, minmetric_7)));

```

The following is a corresponding parameter table.

Parameter	Value
Vector length	2
Reused data set size	-
I/O IRAMs	8 I + 1 O
ALU	8*4*min = 32
BREG	8*(3*shln+3*shrn)= 48
FREG	0
Data flow graph width	8*4= 32

Data flow graph height	5
Configuration cycles	8+2

Re-Normalization

The fourth loop subtracts the minimum of the third loop from each element in the array. The read-modify-write operation has to be broken up into two IRAMs. Otherwise, the IRAM ports will limit throughput.

```

char * iram0= vp->new _metrics; // XPPPreload      (0, vp->new
 _metrics, 64/4)
char *iram1= vp->new _metrics; // XPPPreloadClean(1, vp->new
 _metrics, 64/4)

for (i =0; i<64; i++)
    *iram1++ = *iram0++ - minmetric;

```

The following is a corresponding parameter table.

Parameter	Value
Vector length	64
Reused data set size	-
I/O IRAMs	2 I + 1 O
ALU	1
BREG	0
FREG	0
Data flow graph width	1
Data flow graph height	1
Configuration cycles	64

There are no loop carried dependencies. Since the data size is bytes, the inner loop can be unrolled four times without exceeding the IRAM bandwidth requirements. Networks splitting the 32-bit stream into 4 8-bit streams and rejoining the individual results to a common 32-bit result stream are inserted.

```

char *iram0= vp->new _metrics; // XPPPreload      (0, vp->new
 _metrics, 16)
char *iram1= vp->new _metrics; // XPPPreloadClean      (1, vp-
>new metrics, 16)
for (i=0; i<16; i++) {
    *iram1++ = *iram0++ - minmetric;

```

```

    *iram1++ = *iram0++ - minmetric;
    *iram1++ = *iram0++ - minmetric;
    *iram1++ = *iram0++ - minmetric;
}

```

5

The following is a corresponding parameter table.

Parameter	Value
Vector length	16
Reused data set size	-
110 IRAMs	$2I + 1O$
ALU	$4 \times 4(\text{sub}) = 16$
BREG	$6 \times \text{shln} + 6 \times \text{shrn} = 12$
FREG	0
Data flow graph width	4
Data flow graph height	5
Configuration cycles	$2(\text{split}) + 4 \times 1(\text{sub}) + 2(\text{join}) = 8$

Unroll and jam can be applied after loop tiling, in analogy to the third loop, but loop tiling is now limited by the BREGs used by the split and join networks. The computed tiling size (unroll factor) is $64 \text{ BREGs} / 12 \text{ BREGs} = 5$, which is replaced by 4, since the same throughput is achieved with less over-head.

```

char *iram0= vp->new _metrics;                                //
XPPPreload          (0, vp->new _metrics, 4)
char *iram1= vp->new _metrics ;                                //
15 XPPPreloadClean (1, vp->new _metrics,      4 )
   char *iram2= vp->new _metrics+16;          // XPPPreload      (2,
   vp->new _metrics+16,      4)
   char *iram3= vp->new _metrics+16;          // XPPPreloadClean
                                   (3, vp->new _metrics+16, 4 )
20 char *iram4= vp->new _metrics+32;          // XPPPreload      (4,
   vp->new _metrics+32, 4)
   char *iram5= vp->new _metrics+32;          // XPPPreloadClean
                                   (5, vp->new _metrics+32, 4)
   char *iram6= vp->new _metrics+48;          //
25 XPPPreload          (6, vp-> new _metrics+48, 4)
   char *iram7= vp->new _metrics+48;          // XPPPreloadClean
                                   (7, vp->new _metrics+48, 4)

```

```

for (i=0; i<4; i++) {
*iram1++ = *iram0++ - minmetric;           // first pipeline
    *iram1++ = *iram0++ - minmetric;
5    *iram1++ = *iram0++ - minmetric;
    *iram1++ = *iram0++ - minmetric;
    *iram3++ = *iram2++ - minmetric;       //
second    pipeline
        *iram3++ = *iram2++ - minmetric;
10        *iram3++ = *iram2++ - minmetric;
        *iram3++ = *iram2++ - minmetric;
        *iram5++ = *iram4++ - minmetric;   //
third pipeline
        *iram5++ = *iram4++ - minmetric;
15        *iram5++ = *iram4++ - minmetric;
        *iram5++ = *iram4++ - minmetric;
        *iram7++ = *iram6++ - minmetric;   //
fourth pipeline
        *iram7++ = *iram6++ - minmetric;
20        *iram7++ = *iram6++ - minmetric;
        *iram7++ = *iram6++ - minmetric;
}

```

The following is a corresponding parameter table.

Parameter	Value
Vector length	4
Reused data set size	-
I/O IRAMs	51+40

ALU	$4*(6(\text{split})+4(\text{sub})+6(\text{join})) = 64$
BREG	$4*(6*\text{shln}+6*\text{shrn}) = 48$
FREG	0
Data flow graph width	16
Data flow graph height	1
Configuration cycles	$2(\text{split})+4*1(\text{sub})+2(\text{join}) = 8$

5.4.6—Final Code

Finally ~~we arrive at~~ the following code may be obtained:

```

int update_viterbi29 (void *p, unsigned char sym1, unsigned
5 char sym2) {    int i;
    struct v29 *vp =
p;    unsigned char
*tmp;
    int normalize = 0;
10
    // initialization loop eliminated

    // for (i=0; i<8; i++)
    //    vp->dp->w [i] = 0;

15
    // Configuration for butterfly loop
    char                *iram0= Branchtab29_1;    //
XPPPreload(0, Branchtab29_1, 128/4);
    char    *iram2= Branchtab29 _2;    // XPPPreload (2,
20 Branchtab29_2, 128/4);
    char    *iram4= vp->old_metrics; // XPPPreload(4, vp->old
metrics, 128/4);
    char    *iram5= vp->old _metrics+128;    // XPPPreload(5,
vp->old_metrics+128,128/4);
25    short *iram6= vp->new_metrics;    // XPPPreload (6, vp-
>new_metrics, 128/2);
    unsigned long *iram7 = vp->dp->w; // XPPPreload (7, vp->dp-
>w, 8);

```

```

// sym1 & sym2 are in IRAM 1 & 3
for (_i=0; _i<8; _i++) {
    rlse= 0;
    for (i2=0; i2<32; i2+=2) { // unrolled once
5      unsigned char metric, _tmp, m0, m1, _m0, _m1

        metric = ((*iram0++ ^ sym1) +
                  (*iram1++ ^ sym2) + 1) /2;
        _tmp= (metric << 1) -15;
10

        m0 = *iram2++ + metric;
        m1 = *iram3++ + (15 - metric);
        _m0 = m0 - _tmp;
        _m1 = m1 + _tmp;
15      *iram6++ = (min (m0, m1) << 8) | min (_m0, _m1);
        rlse = rlse | ( m0 >= m1) << i2
                  | (_m0 >= _m1) << (i2+1);
    }
    *iram7++ = rlse;
20 }

/* Renormalize metrics */
if (vp->new _metrics [0] > 150) {
    int 1;
25

// Configuration for loop 3
    char                                *iram0= vp->new _metrics;
// XPPPreload (0, vp->new _metrics, 8);
    char *iram1= vp->new _metrics+8; //    XPPPreload (1, vp-
30 >new _metrics+8, 8);
    char                                *iram2= vp->new _metrics+16;
    //                                XPPPreload (2, vp->new
    _metrics+16, 8);

```



```

        char                                *iram3= vp->new _metrics+24;
        //                                XPPPreload (3, vp->new
_metrics+24, 8);

        char                                *iram4= vp->new _metrics+32;
5      //                                XPPPreload (4, vp->new
_metrics+32, 8);

        char                                *iram5= vp->new _metrics+40;
        //                                XPPPreload (5, vp->new
_metrics+40, 8);
10     char                                *iram6= vp->new _metrics+48;
        //                                XPPPreload (6, vp->new
_metrics+48, 8);

        char                                *iram7= vp->new _metrics+56;
        //                                XPPPreload (7, vp->new
15 _metrics+56, 8);

```

```

for (i=0; _i<2; i++) {
    minmetric0 = min (minmetric0, min( min(*iram0++,
20   *iram0++),
                                min(*iram0++, *iram0++)));

    minmetric1 = min (minmetric1, min( min(*iram1++,
    *iram1++),
                                min(*iram1++, *iram1++)));
25   minmetric2 = min (minmetric2, min( min(*iram2++,
    *iram2++),
                                min(*iram2++, *iram2++)));

    minmetric3 = min (minmetric3, min( min(*iram3++,
    *iram3++),
30   min(*iram3++, *iram3++)));

    minmetric4 = min (minmetric4, min( min(*iram4++,
    *iram4++),
                                min(*iram4++, *iram4++)));

```

```

        minmetric5 = min (minmetric5, min( min(*iram5++,
*iram5++),
                                min(*iram5++, *iram5++)));
        minmetric6 = min (minmetric6, min( min(*iram6++,
5  *iram6++),
                                min(*iram6++, *iram6++)));
        minmetric7 = min (minmetric7, min( min(*iram7++,
*iram7++),
                                min(*iram7++, *iram7++)));
10     }
    minmetric = min(min(( min(minmetric_0, minmetric_1),
                            min(minmetric_2, minmetric_3)),
                    min(( min(minmetric_4, minmetric_5),
                            min(minmetric_6, minmetric_7)));
15 // minmetric is written to the output IRAM

// Configuration for loop 4, minmetric is in an input IRAM
char *iram0= vp->new_metrics;           // XPPPreload
(0, vp->new_metrics, 4)
20 char *iram1= vp->new_metrics;         // XPPPreloadClean
(1, vp->new_metrics, 4)
char *iram2= vp->new_metrics+16;        // XPPPreload
(2, vp->new_metrics+16, 4)
char *iram3= vp->new_metrics+16;        // XPPPreloadClean
25 (3, vp->new_metrics+16, 4)
char *iram4= vp->new_metrics+32;        // XPPPreload
(4, vp->new_metrics+32, 4)
char *iram5= vp->new_metrics+32;        // XPPPreloadClean
(5, vp->new_metrics+32, 4)
30 char *iram6= vp->new_metrics+48;      // XPPPreload
(6, vp->new_metrics+48, 4)
char *iram7= vp->new_metrics+48;        // XPPPreloadClean
(7, vp->new_metrics+48, 4)
    for (i=0; i<4; i++) {

```

```

        *iram1++ = *iram0++ - minmetric;                // first
pipeline
        *iram1++ = *iram0++ - minmetric;
        *iram1++ = *iram0++ - minmetric;
5         *iram1++ = *iram0++ - minmetric;
        *iram3++ = *iram2++ - minmetric;                // second
pipeline
        *iram3++ = *iram2++ - minmetric;
        *iram3++ = *iram2++ - minmetric;
10        *iram3++ = *iram2++ - minmetric;
        *iram5++ = *iram4++ - minmetric;                // third
pipeline
        *iram5++ = *iram4++ - minmetric;
        *iram5++ = *iram4++ - minmetric;
15        *iram5++ = *iram4++ - minmetric;
        *iram7++ = *iram6++ - minmetric;                // fourth
pipeline
        *iram7++ = *iram6++ - minmetric;
        *iram7++ = *iram6++ - minmetric;
20        *iram7++ = *iram6++ - minmetric;
    }
    normalize = minmetric;
}

25    vp->dp++;
    tmp = vp->old_metrics;
    vp->old_metrics = vp->new_metrics;
    vp->new_metrics = tmp;

30    return normalize;
}

```

Performance Considerations

In this example ~~we do~~there is not have a high data locality. Every input data item is read exactly once. Only in the case of re-normalization, the *new_metric* array is ~~re-read~~re-read and re-written. To fully utilize the PAE array, loop tiling was ~~used~~ in conjunction with reduction recognition to break dependencies using algebraic identities. In some cases (minimum search) this ~~leads~~may lead to extremely short vector lengths. This ~~does is not hurt a~~ problem as it still does reduce the running time of the configuration and the transfer time from the top of the memory hierarchy to the IRAMs stays the same. The vector length ~~could~~can be increased if the outer loop that calls the function ~~was is known—the~~ The additional data ~~could~~can be used to increase the fill grade of the IRAMs by unrolling the outer loop, keeping the vector length longer. This would further increase configuration performance by reducing overall pipeline setup times.

Performance of XPP for this example is compared to a hypothetical superscalar RISC-architecture. ~~We assume an~~An average issue width of two ~~is assumed~~, which means that the RISC on average executes two operations in parallel. The estimate is achieved by counting instructions for the source code in 5.42, ~~presented under the heading “Interprocedural Optimizations and Scalar Transformations.”~~ See the table below.

Operation	Cycles	Bfly Setup	Butterfly	Min Setup	Min Search	Norm Setup	Normalize
ADRCOMP	1	8	7		1		
LD/ST	2	5	8	2		1	2
LDI	1	3	4	1		1	
MOVE	1		4		1		
BITOP	1		10				
ADD/SUB	1		20		3	1	3
MULT	2	2	3		2		1
CJMP	3						
Cycles		23	70	5	11	4	10
Count		1	128	1	64	1	64
Issue Width	2						
Total Cycles		12	4480	3	352	2	320

Est. RISC cycles
5168 RISC Cycles

5.5 — MPEG2 encoder/decoder

5.5.1 — Quantization /Inverse Quantization (quant.c)

The quantization file ~~contains~~may include routines for quantization and inverse quantization of 8x8 macro blocks. These functions may differ for intra and non-intra blocks ~~and~~ furthermore. Furthermore, the encoder ~~distinguishes~~may distinguish between MPEG1 and MPEG2 inverse quantization.

This ~~gives~~may give a total of 6 functions, which are all candidates for function inlining, since they do not use the XPP capacity by far.

Since all functions may have the same layout (some checks, one main loop running over the macro block quantizing with a quantization matrix), ~~we concentrate on~~focus is placed on “iquant_intra,” the inverse quantization of intra-blocks, since it ~~contains~~may include all elements found in the other procedures. (The non_intra quantization loop bodies are more complicated, but add no compiler complexity). In the source code the mpeg1 part is already inlined, which is straightforward since the function is statically defined and ~~contains~~includes no function calls itself. Therefore, the compiler inlinesmay inline it and dead function elimination ~~removes~~may remove the whole definition.

Original Code

```
void iquant_intra(src,dst,dc_prec,quant_mat,mquant)
15  short *src, *dst;
    int dc_prec;
    unsigned char *quant_mat;
    int mquant;
    {
20     int i, val, sum;

        if (mpeg1) {
            dst[0] = src[0] << (3-dc_prec);
            for (i=1; i<64; i++)
25         {
            val = (int) (src[i] * quant_mat[i] * mquant) / 16;

            /* mismatch control */
            if- ((val&1)= 0 && val!=0)
30         val+= (val>0) ? -1 : 1;

            /* saturation */
            dst[i] = (val>2047) ? 2047 : ((val<--2048) ? -2048 :
val);
```

```

    }
}
else
{
5      sum = dst[0] = src [0] << (3- dc_prec);
      for (i=1; i<64; i++)
      {
          val = (int) (src [i]*quant_mat [i] *mquant) /16;
          sum+= dst[i] = (val>2047) ? 2047 : ((val<-2048) ? -2048
10 : val);
      }

      /* mismatch control */
      if ((sum&1)==0)
15      dst [63] ^=1;
    }
}

```

Interprocedural Optimizations

20 ~~Analysing~~Analyzing the loop bodies ~~shows, it can be seen~~ that they may easily fit to the XPP and do not use the maximum of resources by far. The function is called three times from module putseq. c. With inter-module function inlining, the code for the function call ~~disappears~~may disappear and ~~is~~may be replaced with the function . Therefore, it readsmay be as follows:

```

25  for (k=0; k<mb _height*mb _width; k++) {
      if (mbinfo [k] . mb _type & MB_INTRA)
          for (j =0; j<block _count; j ++)
              if (mpeg1) {

30          blocks[k*block_count+j] [0] = blocks[k*block_count+j] [0]
              <<
                  (3-dc_prec);

          for (i=1; i<64; i++) {

```

```

        val = (int) ( blocks [ k*block_count+j] [i]  * intra_q
[i] *mquant) /16;

        . . .

    }

5    }else {

        sum = blocks[k*block_count+j][0] =
blocks[k*block_count+j][0] <<

                                (3-dc_prec);

        for (i=1; i<64; i++) {
10        val = (int)(blocks[k*block_count+j][i] *
intra_q[i]*mquant) /16;

            . . .

        }
        }else {
15        . . .

        }

```

Basic transformations

Since global mpeg 1 does not change within the loop, unswitching ~~moves~~may move the

20 control statement outside the j loop and ~~produces~~may produce two loop nests.

```

for (k=0; k<mb_height*mb_width; k++) {
    if (mbinfo[k].  mb_type & MB_INTRA)
        if (mpeg1)
            for (j=0; j<block _count; j++) {
25        blocks[k*block_count+j][0] = blocks[k*block_count+j][0]
<<                                (3-dc_prec);

                for (i=1; i<64; i++) {
                    val = (int) ( blocks [k*block _count+j] [i] * intra_q[i]
                    *mquant) /16;
30        . . .

                }

```

```

    }
    else
        for (j=0; j<block_count; j++) {
            sum = blocks [ k*block_count+j ] [0] =
5   blocks[k*block_count+j][0] <<
                                   (3-dc_prec);

            for (i=1; i<64; i++) {
                val = (int) ( blocks [ k*block_count+ j ] [i ] *
intra_q [i] *mquant) /16;
10         . . .
            }
        }
    }
}

```

15 Furthermore, the following transformations ~~are done~~may be performed:

- A peephole optimization ~~reduces~~may reduce the divide by 16 to a right shift 4. This is~~may be~~ essential since ~~we do not consider~~ loop bodies ~~containing~~including division for the XPP are not considered.
- Idiom recognition ~~reduces~~may reduce the statement after the “saturation” comment to
20 dst [i] = min(max(val, -2048), 2047);

Increasing parallelism

~~Now we want~~It may be desired to increase parallelism. The j-i loop nest is a candidate for unroll-and-jam when the interprocedural value range analysis finds ~~out~~ that block_count can
25 only get the values 6, 8, or 12. Therefore, it has a value range [6,12] with the additional attribute to be dividable by 2. Thus, an unroll_and_jam with the factor 2 is applicable (the resource constraints would choose a ~~bigger~~greater value). Since no loop carried dependencies exist, this transformation is safe.

30 ~~It~~This is to say that the source code contains a manually peeled first iteration. This peeling has been done because the value calculated for the first block value is completely different from the other iterations, and the control statement in the loop would cause a major performance decrease on traditional processors. Although this does not prevent unroll-and-

jam (because there are no dependencies between the peeled ~~off~~ first iteration and the rest of the loop), the transformation must be prepared to handle such cases.

After unroll ~~and~~ jam, the source code ~~looks like~~ may be approximately as follows (only one

5 of the nests ~~showed~~shown and the peeled first iterations moved in front):

```
for (j=0; j<block _count; j+=2) {
    blocks [k*count+j ] [0] = blocks [k*count+j ] [0] << (3-dc
prec);    blocks [k*count+j+1] [0] = blocks [k*count+j+1] [0]
<< (3-dc prec) ;
```

```
10 for (i=1; i<64; i++) {
    val = (int) (blocks[k*count+j][i]*intra_q[i] *mbinfo [k] .
mquant) >>4;
```

```
/* mismatch control */
```

```
15 if ((val &1) ==0 && val ! =0 )
```

```
    val+= (val>0) ? -1 : 1;
```

```
/* saturation */
```

```
blocks[k*count+j][i] = min(max(val, -2048), 2047);
```

20

```
    val =
(int) (blocks[k*count+j+1][i]*intra_q[i]*mbinfo[k].mquant) >>4;
```

```
/* mismatch control */
```

```
25 if ((val&1)==0 && val!=0)
```

```
    val+= (val>0) ? -1 : 1;
```

```
/* saturation */
```

```

        blocks[k*count+j+1][i] = min(max(val, -2048) , 2047) ;
    }
}

```

5 Further parallelism can be obtained by index set splitting. Normally used to break dependence cycles in the DDG, it can here be used to split the i-loop in two and let two sub-
~~on configurations~~⁵ configurations (sub-configuration is chosen as a working title for
configurations that include independent networks that do not interfere) work on distinct
blocks of data. Thus, the i loop is split into 2 or more loops which work on different subsets
10 of the data at the same time.

Handling the data types

In contrast to the FIR-Filter, edge detector and matrix multiplication benchmarks, which all
use data types fitting perfectly to the XPP^[6] (it is assumed that the size of int is chosen to
15 be the XPP architecture data bit width, as everything else would not lead to any feasible
result), the MPEG2 codec uses all data types commonly used on a processor for desktop
applications. Written for the Intel x86 and comparable architectures, ~~we must assume~~ it may
be assumed that the sizes of char, short, and int are 8,16, and 32 respectively. Assuming that
the XPP has a bit width of 32 ~~we must take precautions~~ should be taken for the smaller data
20 types.

Therefore ~~we split~~, the stream of data packets with each packet ~~containing~~ including 2 or 4
values of the shorter data type may be split into 2 or 4 streams. If ~~we have~~ enough resources
are left, this will cause no performance penalty. Each of the divided streams ~~is~~ may be sent to
25 its own calculation network; ~~therefore~~. Therefore, in every cycle, two short or four char
values ~~are~~ may be handled. Nevertheless, this ~~causes~~ may cause an area penalty; because,
besides the split-merge elements, the whole data flow graph has to be duplicated as often as
needed. ~~Figure 64~~ Fig. 41 shows how short values are handled. It shows the splitting of short
values into two streams and the merging of the streams after the calculation. The packet is

⁵ ~~sub-configuration is chosen as a working title for configurations which contains independent networks that do not interfere:~~

⁶ ~~We assume that the size of int is chosen to be the XPP architecture data bit width. Everything else would not lead to any feasible result~~

split into its *hi*- and *lo* part by shift operations and merged behind the calculation branches. The legality of this transformation is the same as with loop unrolling, with an unrolling factor as big as the data type ~~is~~being smaller as the architecture data type.

5 ~~Unfortunately this~~This, however, is not the end of the pole. ~~The~~It may be further required for the compiler ~~further has to assure~~ensure that every intermediate result which produces an over/under-flow for the shorter data type does the same with the bigger data type. Therefore, it has to insert clipping operations which ~~assure~~ensure that the network calculates with real 16 or 8 bit ~~value~~values, respectively.

10

~~Figure 64 Splitting short values into two streams and merging them
after the calculation. This method causes no performance penalty~~

If the configuration size does not allow the whole loop body to be duplicated or dependencies prevent this, ~~we~~there is still have thea possibility to ~~merge~~of merging the split values again.

15 ~~This of course~~ causes a performance penalty to the previous solution, because the throughput is only one (short) value/cycle ~~now~~. ~~Figure 65~~Fig. 42 shows how the merge is done. Instead of streaming parallel through two networks, the values are serialized and de-serialized again after the network. ~~Figure 65 Merging the~~The split values are merged before the network. An event ~~gen-erator~~generator drives the merge and ~~demux~~Demux PAEs. ~~This figure~~Fig. 42
20 replaces the ~~2 black~~two boxes labeled “network” in ~~Figure 64~~Fig. 41.

5.5.2—Inverse Discrete Cosine Transformation (idct.c)

The idct-algorithm ~~is~~may be used for the MPEG2 video decompression algorithm. It operates on 8x8 blocks of video images in their frequency representation and transforms them
25 back into their original signal form. The MPEG2 decoder contains a transform-function that calls idct for all blocks of a frequency-transformed picture to restore the original image.

The idct function ~~consists of~~may include two for-loops. The first loop calls idctrow, and the second calls idctcol. Function inlining is able to eliminate the function calls within the entire
30 loop nest structure so that the numeric code is not interrupted by function calls anymore. ~~Another~~In another embodiment, a way to get rid of function calls between the loop nest is loop embedding that pushes loops from the caller into the callee.

Original Code (idct.c)

```
/* two dimensional inverse discrete cosine transform */
void idct (block)
  short *block;
5  {
    int i;

    for (i=0; i<8; i++)
      idctrow (block+8 *i)
10    ;

    for (i=0; i<8; i++)
      idctcol (block+i)
      ;
15  }
```

The first loop ~~changes~~may change the values of the block row by row. Afterwards, the changed block is further transformed column by column. ~~All~~In this embodiment, all rows have to be finished before any column processing can be started. The function is illustrated
20 in Fig. 43.

Dependency analysis ~~detects~~may detect true data dependencies between row processing and column processing. Therefore, it may be required for the processing of the columns ~~has to be~~ delayed until all rows are done. The innermost loop bodies idctrow and idctcol are nearly
25 identical. They process numeric calculations on eight input values (column values in case of idctcol and row values in case of idctrow). Eight output values are calculated and written back (as column/row). Idctcol additionally applies clipping before the values are written back. ~~This is why we concentrate on idctcol~~Accordingly, idctcol is presented herein. The
code may be as follows:

```
30 /* column (vertical) IDCT
   *
```

```

*                               7                               pi
* dst[8*k] = sum c [1] * src [8*1] * cos(-- *
*                               1=0                               8

(k +  $\frac{1}{2}$ ) *1)

*
* where: c[0] = 1/1024
5 * c[1..7] = (1/1024)* sqrt(2)
*/
static void idctcol (blk)
short *blk;
{
10 int x0, x1, x2, x3, x4, x5, x6, x7, x8;

/* shortcut */
if (! (x1 = (blk[8*4]<<8))) | (x2 = blk[8*6]) |

(x3 = blk[8*2]) | (x4 = blk[8*1]) | (x5 = blk[8*7]) |
15 (x6 = blk[8*5]) | (x7 = blk[8*3]))
{
blk[8*0]=blk[8*1]=blk[8*2]=blk[8*3]=blk[8*4]=blk[8*5] =
blk[8*6]=blk[8*7]= iclp[ (blk[8*0] +32)>>6];
return;
20 }

x0 = (blk [8*0] <<8) + 8192;

/* first stage */
x8 = W7* (x4+x5) + 4;
25 x4 = (x8+ (W1-W7) *x4) >>3;

x5 = (x8- (W1+W7) *x5)>>3;
x8 = W3* (x6+x7) + 4;
x6 = (x8- (W3-W5) *x6)>>3;

```

```

x7 = (x8 - (W3+W5)*x7)>>3;

/* second stage */
x8 = x0 + x1;
5  x0 -= x1;
x1 = W6* (x3+x2) + 4 ;

x2 = (x1 - (W2+W6)*x2)>>3
x3 = (x1+(W2-W6)*x3)>>3;

10  x1 = x4 + x6;

x4 -= x6;

x6 = x5 + x7;

x5 -= x7;

15  /* third stage * /
x7 = x8 + x3;
x8 -= x3;
x3 = x0 + x2;
x0 -= x2;
20  x2 = (181*(x4+x5)+128)>>8;
x4 = (181*(x4-x5)+128)>>8;

/* fourth stage * /
blk[8*0] = iclp[(x7+x1)>>14];
25  blk[8*1] = iclp[(x3+x2)>>14];
blk[8*2] = iclp[(x0+x4)>>14];
blk[8*3] = iclp[(x8+x6)>>14];
blk[8*4] = iclp[(x8-x6)>>14];
blk[8*5] = iclp[(x0-x4) >>14 ] ;
30  b1 k [ 8 * 6 ] = iclp [ (x3-x2) >>14 ] ;
blk[8*7] = iclp[ (x7-x1)>>14];

```

```
}
```

W1 - W7 are macros for numeric constants that are substituted by the preprocessor. The iclp array is used for clipping the results to 8-bit values. It is fully defined by the init_idct

5 function before idct is called the first time:

```
void init_idct()
{
    int i;

10    iclp = iclip+512;
    for (i= -512; i<512; i++)
        iclp[i] = (i<-256) ? -256 : ((i>255) ? 255 : i) ;
}
```

15 A special kind of idiom recognition (function recognition) is able to replace the calculation of each array element by a compiler known function that can be realized efficiently on the XPP. If the compiler features whole program memory aliasing analysis, it is able to replace all uses of the iclp array with the call of the compiler known function. Alternatively, a developer can replace the iclp array accesses manually by the compiler known saturation function calls. The illustration. Fig. 44 shows a possible implementation for saturate(val,n) as an NML schematic using two ALUs. In this case, it is necessary to replace array accesses like iclp[i] by with saturate(i,256).

25 The `/*shortcut*/` code in idctcol ~~speeds~~may speed column processing up if x1 to x7 is zero. This breaks the well-formed structure of the loop nest. The if-condition is not loop invariant and loop unswitching cannot be applied. ~~But nonetheless~~ Nonetheless, the code after shortcut handling is well suited for the XPP. It is possible to synthesize if-conditions for the XPP (speculative processing of both blocks plus selection based on condition) but this would just waste PAEs without any performance benefit. Therefore, the `/*shortcut*/` code in idctrow and idctcol has to be removed manually. The code snippet below shows the inlined version of the idctrow-loop with additional cache instructions for XPP control:

```
void idct (block)
short *block;
```

```

{
    int i;

    XPPPreload(IDCTROW _CONFIG); // Loop Invariant
5
    for (i=0; i<8; i++) {
        short *blk;
        int x0, x1, x2, x3, x4, x5, x6, x7, x8;
        blk = block+8*i;
10
        XPPPreload (0, blk, 8) ;
        XPPPreloadClean(1, blk, 8) ; // IRAM1 is erased and
        assigned to blk
        XPPExcute(IDCTROW_CONFIG, IRAM(0) ; IRAM(1)) ;
15    }
    for (i=0; i<8; i++) {
        . . .
    }
}

```

20

As the configuration of the XPP does not change during the loop execution, invariant code motion has moved out **XPPPreload**(IDCTROW_CONFIG) from the loop.

NML Code Generation

25

Data Flow Graph

As idctcol is more complex due to clipping at the end of the calculations ~~we decided to take~~, idctcol is well suited as a representative loop body for a presentation of the data flow graph.

~~The figure on the next page~~ Fig. 45 shows the data flow graph for the
 30 IDCTCOLUMN_CONFIG. A heuristic has to be applied to the graph to estimate the resource needs on the XPP. In ~~our~~ this example, the heuristic produces the following results:

	ADD, SUB	MUL	<<X, >>X	Saturate (x,n)
Ops needed	35	11	18	8
	ALUs	FREGs	BREGs	
Res. left	19	80	45	
Res. avail.	64	80	80	

~~Fortunately the~~ The data flow graph fits into an XPP64 and ~~we can~~ this example may proceed without *loop dissection*^[7] (splitting the loop body into suitable chunks) ~~for this example.~~

5 See João M.P. Cardoso et al., *supra*.

Address Generation:

To fully synthesize the loop body ~~we have to face~~ the problem of address generation for accessing the data must be addressed.

10

For IDCTCOLUMN_CONFIG ~~we have to select~~, the n^{th} element of every row must be selected, which means an address serial of (0,8,16. . . 1,9,17. . . 7,15,23. . .). ~~We use two~~ Two counter macros may be used for address generation as shown ~~opposite~~ in Fig. 46.

15 The upper counter increments by eight and the lower counter increments by one. The IRAM output is passed to the data flow graph of IDCTCOLUMN. If all (eight) row elements of a column are available, SWAP is switched through to the data flow graph input and the calculation for a new column begins.

20 For the IDCTROW_CONFIG, the address generation is very simple as the IRAM already ~~contains~~ has the block in the appropriate order (row after row as it has to be accessed). Again, by using SIUP (stepped iterative up)-counter macros as described in the XPP tutorial, it is possible to map linear address expressions to NML-code in a generic way. As IDCTROW_CONFIG accesses a two-dimensional array ~~we need~~, two SIUP-counters may be needed in the corresponding NML code. The column-elements have to be accessed row after
25 row so the upper ~~counter~~ counter's increment is one and the lower ~~counter~~ counter's increment is eight. However, the NML code for this access pattern (0. . . 5,6,7,8,9. . . 63) can be reduced to one single counter (or to FIFO-mode IRAM access).

⁷ ~~XPP VC: A C Compiler with Temporal Partitioning for the PACT-XPP Architecture, J. M. P. Cardoso and Markus Weinhardt~~

Address generation for write access ~~is~~may be implemented in the same manner. The resources have to be updated to take this additional code into account. It takes $2 \cdot (8+8+2 \cdot 1)$ FREGs and $2 \cdot (2+1)$ more BREGs in the worst case, which is still available on the XPP.

5

If IRAM use is not critical, it is also possible to distribute the data on several IRAMs to improve the memory throughput into the XPP-array. This task ~~has to~~may be done by the RISC-core or by a more sophisticated XPP-cache controller.

10 **Further Enhancing XPP Utilization**

As mentioned ~~at the beginning above~~, idct is called for all data blocks of a video image (loop in transform. c). This circumstance ~~allows us to further improve~~may allow for improvement of the XPP utilization.

15 When ~~we look~~looking at the data flow graph of idctcol in detail ~~we see~~, it can be seen that it forms a very deep pipeline. ~~If we bring back to our mind~~Considering that the IDCTROW_CONFIG runs only eight times on the XPP, which ~~meant~~means that only 64 (8 times 8 elements of a column) elements are processed through this pipeline, and that ~~we have to wait then until all data left the pipeline before we can change~~ from the XPP configuration to the IDCTCOLUMN_CONFIG configuration to go on with column processing ~~then it gets obvious that something must wait until all data has left the pipeline, this example is suboptimal in our example.~~

20

Problem (Pipeline- Depth)

25 The pipeline is just too deep for processing only eight times eight rows. Filling and flushing a deep pipeline is expensive if only little data is processed with it. First the units at the end of the pipeline are idle and then the units at the ~~begin~~beginning are unused—, as shown in Fig. 47.

30 **Solution (Loop Tiling)**

It is profitable to use loop interchange for moving the dependencies between row- and column processing to an outer level of the loop nest. The loop that calls the idct-function (in transform. c) on several blocks of the image ~~is~~ has no loop interchange preventing

dependencies. Therefore, this loop can be moved inside the loops of column and row processing, as shown in Fig. 48.

Now the processing of rows and columns can be applied on more data (by applying loop tiling) ~~and therefore.~~ Therefore, filling and flushing the pipeline can be neglected.

Constraints (Cache Sensitive Loop Tiling)

The caching hierarchy has to be taken into account when ~~we defined~~ defining the number of blocks that will be processed by the IDCTROW_CONFIG. ~~Remember, we need~~ As discussed above, the same blocks are needed in the subsequent IDCTCOLUMN_CONFIG configuration! ~~We have to take care.~~ It should be ensured that all blocks that are processed during IDCTROW_CONFIG fit into the cache. Loop tiling has to be applied with respect to the cache size so that the- processed data fits into the cache.

IRAM reuse between different configurations

This example implies another bandwidth optimization that is just ~~a more consequent~~ another version of loop tiling. Instead of transferring data from row processing to column processing via the memory hierarchy (cache sensitive loop tiling takes care that only the cache memory is accessed) ~~we can completely bypass,~~ the memory interface can be completely bypassed by using the output IRAM of Config A as input IRAM of Config B, as shown in Fig. 49.

Putting all together

If we apply cache sensitive loop tiling, IRAM reuse, and function in-lining ~~we,~~ the example can be further optimize our example: optimized.

Finally, the idct-function ~~gets~~ becomes completely inlined in transform.c. If block_count is, e.g., 6 and ~~we assume~~ it is assumed that 64*6 words do not exceed the cache size, then ~~we can transform the example~~ may be transformed to:

```
// transform.  c
..
block = blocks [ k* 6 ] ;
XPPPreload ( IDCTROW _CONFIG) ;
```

```

XPPPreload(0,block,64*6) ; //IRAM0 gets 64 words from 6
blocks XPPPreloadClean (1, block, 64*6) ; // erase IRAM1 and
assign to the 6 blocks
XPPExecute (IDCTROW _CONFIG, IRAM(0) , IRAM (1)) ;

```

5

```

XPPPreload ( IDCOLUMN _CONFIG) ;
XPPPreload (1,block,64 *6 ) ; // redundant -> will be
eliminated XPPExecute(IDCOLUMN_CONFIG, IRAM(1), IRAM(2));
..

```

10

The address generation in IDCTROW_CONFIG and IDCOLUMN_CONFIG has to be modified for reflecting the different data block size - caused by loop tiling - that has to be processed. This can be implemented by an additional SUIP counter that generates the block offsets inside the tiles—, as shown in Fig. 50.

15

The following table ~~contains~~provides architectural parameters for IDCTROW_CONFIG and IDCOLUMN_CONFIG of the final result. It relies on a cache that is able to store block _count blocks. As two configurations are executed in this example, the configuration cycles have to be taken twice ~~and therefore~~. Therefore, the total configuration cycles are $2 \times (\text{block_count} \times 64 + (12 + 2 \times 8) \times 2)$.

20

Parameter	Value
Vector length	8 words
Reused data set size	block_count x 64 words
I/O IRAMs	3 (one shared)
ALU	45 FUs
DREG	41 FUs
FREG	36 FUs
Data flow graph width	8
Data flow graph height	12
Configuration cycles	block_count x 64 + (12 + 2*8) x 2

Performance Considerations

In this example, it is possible to exploit high data locality, which means that many operations are performed on a limited memory range. The performance of the ~~proposed~~ XPP solution of this embodiment is compared to a hypothetical superscalar RISC-architecture. ~~We assume~~

25

an An issue width of two is assumed, which means that the RISC executes on average two operations in parallel.

	Ops for Row/Column	Est. RISC cycles	
LD/ST	16	2	32
ADRCOMP	16	1	16
ADD/SUB	35	1	35
MULT	11	2	22
SHIFT	18	1	18
SAT	8	4	32
Issue Width		2	155
		Cyc/Row(Col)	<u>78</u>
Proc. Rows	8	620	
Proc. Cols	8	620	
RISC Cyc/Blk		<u>1240</u>	
XPP Cyc/Blk		<u>128</u>	with data duplication+reordering 24
Speedup		<u>10</u>	with data duplication+reordering <u>52</u>

5 Even though speedup is reasonable ~~it gets obvious that~~, fetching the input data from a single IRAM (which means that ~~we have~~ it is required to feed the eight inputs in eight cycles before processing is started) reduces the potential speedup significantly. ~~With~~In other words ~~we have~~, there is a pipeline that is able to process eight input values per cycle, but ~~we are loading~~ the pipeline is loaded only every eighth cycle. This causes that only every eighth pipeline stage is filled. ~~The figure below~~Fig. 51 illustrates this:

Full utilization can be achieved only by loading the eight input values of the pipeline in one cycle. A ~~simple~~ solution to improve the memory throughput to the pipeline is data duplication as described ~~in~~under the hardware section. ~~heading "Hardware."~~

15 Instead of loading the six 8x8 blocks to a single IRAM ~~we use~~, in an embodiment of the present invention, the **XPPPreloadMultiple** command may be used to load the eight IRAMs with the same contents:

```
XPPPreload(0,block, 64*6) ; //IRAM0 gets 64 words from
```

20 6 blocks

is changed to:

```
XPPPreloadMultiple(0xFF, block, 64x6) // load RAM0 up to IRAM7
with blocks
```

Now the pipeline gets fully utilized and ~~we also have to store eight results per cycle~~ must be stored. This can be achieved by writing every output value- to another IRAM_i which additionally takes eight more IRAMs (~~using~~ Using data duplication in this example ~~needs~~requires all 16 IRAMs of the XPP64.) For storing the data that is generated with

5 IDCTROW_CONFIG we have to change:

```
XPPPreloadClean (1, block, 64 * 6) ; / / erase IRAM1 and
assign to the 6 blocks
```

to:

```
tmpsize = 64*6/8;
10 XPPPreloadClean(8, block+0*tmpsize, tmpsize); //IRAM8 for
interm. Rslt 1

XPPPreloadClean(9, block+1*tmpsize, tmpsize); //IRAM9 for
interm. Rslt 1

XPPPreloadClean(10, block+2*tmpsize, tmpsize); //      IRAM10
15   for   interm. Rslt 1

XPPPreloadClean(11, block+3*tmpsize, tmpsize); //IRAM11 for
interm. Rslt 1

XPPPreloadClean(12, block+4*tmpsize, tmpsize); //      IRAM12
   for   interm.      Rslt 1

20 XPPPreloadClean(13, block+5*tmpsize, tmpsize); //      IRAM13
   for interm. Rslt 1

XPPPreloadClean(14, block+6*tmpsize, tmpsize); //IRAM14 for
interm. Rslt 1

XPPPreloadClean(15, block+7*tmpsize, tmpsize); //IRAM15 for
25   interm. Rslt 1
```

This causes different data layouts for the intermediate results. ~~We need an~~An additional configuration (REORDER_CONFIG), as shown in Fig. 52, may be needed to restore the original data layout.

30

Again, address generation has to be modified to fetch eight input values per cycle. This, on the one hand, requires seven additional adders, but, on the other hand, avoids swaps and latches for keeping the data eight cycles.

5 Data duplication and data reordering may finally transforms the example code to:

```
// transform c
. .
block = blocks [ k* 6 ] ;
XPPPreload (IDCTROW_CONFIG) ;
10 XPPPreloadMultiple (0xFF, block, 64x6 ) // load IRAM0 up to
    IRAM7 with blocks
    tmpsize = 64 * 6/8; // result gets seperated into 8 IRAMs
XPPPreloadClean(8, block+0*tmpsize, tmpsize); // IRAM 8
    tmpsize); for interm. Rslt 1
15 XPPPreloadClean(9, block+1*tmpsize, tmpsize); // IRAM 9
    tmpsize); for interm. Rslt 1
XPPPreloadClean(10, block+2*tmpsize, tmpsize); // IRAM 10
    tmpsize); for interm. Rslt 1
XPPPreloadClean(11, block+3*tmpsize, tmpsize); // IRAM 11
20 tmpsize); for interm. Rslt 1
XPPPreloadClean(12, block+4*tmpsize, tmpsize); // IRAM 12
    tmpsize); for interm. Rslt 1
XPPPreloadClean(13, block+5*tmpsize, tmpsize); // IRAM 13
    tmpsize); for interm. Rslt 1
25 XPPPreloadClean(14, block+6*tmpsize, tmpsize); // IRAM 14
    tmpsize); for interm. Rslt 1
XPPPreloadClean(15, block+7*tmpsize, tmpsize); // IRAM 15
    tmpsize); for interm. Rslt 1
XPPEXecute(IDCTROW_CONFIG, IRAM(0-7), IRAM(8-15));
30
XPPPreload (IDCOLUMN_CONFIG);
```

```

XPPPpreloadMultiple (0xFF, block, 64x6) / / ld IRAM0-IRAM7 with
interm. Rslt 1
XPPPpreloadClean(8, block+0 *tmpsize, tmpsize); // IRAM8 for
interm. Rslt 2
5 XPPPpreloadClean(9, block+1 *tmpsize, tmpsize); // IRAM9 for
interm. Rslt 2
XPPPpreloadClean(10, block+2 *tmpsize, tmpsize); // IRAM10 for
interm. Rslt 2
XPPPpreloadClean(11, block+3 *tmpsize, tmpsize); // IRAM11 for
10 interm. Rslt 2
XPPPpreloadClean(12, block+4 *tmpsize, tmpsize); // IRAM12 for
interm. Rslt 2
XPPPpreloadClean(13, block+5 *tmpsize, tmpsize); // IRAM13 for
interm. Rslt 2
15 XPPPpreloadClean(14, block+6 *tmpsize, tmpsize); // IRAM14 for
interm. Rslt 2
XPPPpreloadClean(15, block+7 *tmpsize, tmpsize); // IRAM15 for
interm. Rslt 2
XPPEexecute (IDCOLUMN_CONFIG, IRAM(0-7), IRAM(8-15));
20 XPPPpreload ( REORDER_CONFIG) ;
XPPPpreloadMultiple (0xFF, block, 64x6) / /ld IRAM0-IRAM7 with
interm. Rslt 2
rsltsize = 64; // 64*6/6;
XPPPpreloadClean(8, block+0 *rsltsize, rsltsize); // IRAM8 for
25 final Rslt
XPPPpreloadClean(9, block+1 *rsltsize, rsltsize); // IRAM9 for
final Rslt
XPPPpreloadClean(10, block+2 *rsltsize, rsltsize); // IRAM10
for final Rslt
30 XPPPpreloadClean(11, block+3 *rsltsize, rsltsize); // IRAM11
for final Rslt
XPPPpreloadClean(12, block+4 *rsltsize, rsltsize); // IRAM12
for final Rslt

```



```

XPPreloadClean(13, block+4 *rsltsize, rsltsize); // IRAM13
for final Rslt
XPPEecute(IDCOLUMN_CONFIG, IRAM(0-7), IRAM(8-13));

```

5 5.6—Wavelet

5.6.1—Original Code

```

void forward _wavelet ()
{
    int i, nt, *dmid;
10    int * sp , *dp, d_tmp0, d_tmpl, d_tmpl, s_tmp0, s_tmpl;
    int mid, ii;
    int *x;
    int s [256] , d[256] ;

15    for (nt=COL; nt>=BLOCK_SIZE; nt>>=1) {
        for (i=-0; i <nt *COL /* tmp _nt * /; i+=COL) {
            x = & int_data [i] ;
            mid= (nt>>1) -1;

20            s[0] = x[0];

            d[0] = x[ROW];
            s[1] = x[2];
            s[mid] = x[2*mid];

25            d[mid] = x[2*mid+ROW];

            d[0] = (d[0]<<1)-s[0]-s[1];
            s[0] = s[0]+(d[0]>>2);

            d_tmp0 = d[0];
30            s_tmp0 = s[1];

            for (ii=1; ii<mid; ii++)

```

```

    s_tmp1 = x[2*ii+2];
    d_tmp1 = ((x[2*ii+ROW])<<1) - s_tmp0 - s_tmp1;
    d[ii] = dtmp1;
    s[ii]= s_tmp0+((d_tmp0+d_tmp1)>>3);
5    d_tmp0 = d_tmp1;
    s_tmp0 = s_tmp1;
}

```

```

    d[mid] = (d[mid]-s[mid])<<1;

10    s[mid] = s[mid]+((d[mid-1]+d[mid])>>3);

```

```

    for (ii=0; ii<=mid; ii==) {
        x[ii] = s[ii];
        x[ii+mid+1] = d[ii];
15    }
}
for (i=0; i<nt; i++) {

```

```

    x = &int_data[i];
20    mid = (nt>>1)-1;

```

```

    s[0] = x[0];
    d[0] = x[COL];
    s[1] = x[COL<<1];
25    s[mid] = x[(COL<<1)*mid];
    d[mid] = x[(COL<<1)*mid +COL];
    d[0] = (d[0]<<1)-s[0]-s[1];
    s[0] = s[0]+(d[0]>>2);

```

```

30    d_tmp0 = d[0];
    s_tmp0 = s[1];

```

```

    for (ii=1; ii<mid; ii++) {
        s_tmp1 = x[2*COL*(ii+1)];
        d_tmp1 = (x[2*COL*ii+COL]<<1)-s_tmp0-s_tmp1;
        d[ii] = d_tmp1;
5       s[ii] = s_tmp0+((d_tmp0+d_tmp1)>>3);
        d_tmp0 = d_tmp1;
        s_tmp0 = s_tmp1;
    }

10    d[mid] = (d[mid]<<1)-(s[mid]<<1);
    s[mid] = s[mid]+((d[mid-1]+d[mid])>>3);

    for (ii=0; ii<=mid; ii++) {
        x[ii*COL] = s[ii];
15    x[(ii+mid+1)*COL] = d[ii];
    }
}
}
}
}
20

```

5.6.2—Optimizing the Whole Loop Nest

After pre-processing and application of copy propagation over *s_tmp1*, *d_tmp1*, ~~we obtain~~ the following loop nest may be obtained:

```

25 void forward _wavelet ( )
    {
        int i, nt, *dmid;
        int * sp, *dp, d_ tmp0, d_ tmp1, d_ tmpi, s_ tmp0, s_ tmp1;
        int mid, ii;
30    int *x;
        int  s[256], d[256] ;

        for (nt=64; nt>= 16; nt>>=1) {
            for (i=0; i<nt*64; i+=64) {

```

```

x = &int_data[i];
mid = (nt>>1)-1;

5   s[0] = x[0];
    d[0] = x[1];
    s[1] = x[2];
    s[mid] = x[2*mid];
    d[mid] = x[2*mid+1];

10  d[0] = (d[0]<<1)-s[0]-s[1];
    s[0] = s[0]+(d[0]>>2);

    d_tmp0 = d[0];
15  s_tmp0 = s[1];

    for (ii=1; ii<mid; ii++) {
        d[ii] = ((x[2*ii+1])<<1) - s_tmp0 - x[2*ii+2];
        s[ii] = s_tmp0 + ((d_tmp0 + d[ii])>>3);
20  d_tmp0 = d[ii];
    s_tmp0 = s[ii];
    }
    d[mid] = (d[mid]-s[mid])<<1;
    s[mid] = s[mid]+((d[mid-1]+d[mid])>>3);

25  for (ii=0; ii<=mid; ii++) {
        x[ii] = s[ii];
        x[ii+mid+1] = d[ii];
    }
30  }

    for (i=0; i<nt; i++) {
        x = &int_data[i];

```

```

mid = (nt>>1)-1;

s[0] = x[0];

d[0] = x[64];

5   s[1] = x[128];

s[mid] = x[128*mid];

d[mid] = x[128*mid+64];

d[0] = (d[0]<<1)-s[0]-s[1];

s[0] = s[0]+(d[0]>>2);

10

d_tmp0 = d[0];

s_tmp0 = s[1];

for (ii=1; ii<mid; ii++) {
15   d[ii] = (x[128*ii+64]<<1) - s_tmp0 - x[128*(ii+1)];
      s[ii] = s_tmp0 + ((d_tmp0 + d[ii])>>3);
      d_tmp0 = d[ii];
      s_tmp0 = s[ii];
}

20

d[mid] = (d[mid]<<1) - (s[mid]<<1);
s[mid] = s[mid] + ((d[mid-1]+d[mid])>>3);

for (ii=0; ii<=mid; ii++) {
25   x[ii*64] = s[ii];
      x[(ii+mid+1)*64] = d[ii];
}
}

```

```

}
}

```

5 Then we have 4 tables, one Below is a table for each innermost loop. The tables for the first and the third loops are identical, as are the tables for the second and the fourth loop. We have the following two loops. Accordingly, 2 tables are presented below.

Parameter	Value
Vector length	mid-2
Reused data set size	-
I/O IRAMs	6
ALU	6
BREG	0
FREG	2
Data flow graph width	2
Data flow graph height	6
Configuration cycles	6+(mid-2)

Parameter	Value
Vector length	mid
Reused data set size	-
I/O IRAMs	6
ALU	0
BREG	0
FREG	0
Data flow graph width	2
Data flow graph height	1
Configuration cycles	mid

10 The two inner loops do not have the same iteration range and could be ~~candidate~~candidates for loop fusion. ~~Therefore, therefore~~ the first and last iterations of the second loop ~~are~~may be peeled off. The surrounding code between the 2 loops can be moved to after the second loop. Accordingly, then we obtain the following code for the loop nest may be obtained.

```

15 for (nt=64; nt>= 16; nt>>=1) {
    for (i=0; i<nt*64; i+=64) {

```

```

x = &int_data[i];
mid = (nt>> 1)-1;

s[0] = x[0];
5 d[0] = x[1];
s[1] = x[2];

s[mid] = x[2*mid];
d[mid] = x[2*mid+1];
10

d[0] = (d[0]<<1)-s[0]-s[1];
s[0] = s[0]+(d[0]>>2);

d_tmp0 = d[0]
15 s_tmp0 = s[1];

for (ii=1; ii<mid; ii++) {
    d[ii] = ((x[2*ii+1])<<1) - s_tmp0 - x[2*ii+2];
    s[ii] = s_tmp0+((d_tmp0 + d[ii])>>3);
20 d_tmp0 = d[ii];
s_tmp0 = s[ii];
}
for (ii=1; ii<mid; ii++) (
    x[ii] = s[ii];
25 x[ii+mid+1] = d[ii];
}

d[mid] = (d[mid]-s[mid])<<1;

s[mid] = s[mid] + ((d[mid-1]+d[mid])>>3);
30

x[0] = s[0];

```

```

    x[mid+1] = d[0];

    x[mid] = s[mid];

    x[2*mid+1] = d[mid];

}

5
for (i=0; i<nt; i++) {
    x = &int_data[i];
    mid = (nt>>1)-1;

10    s[0] = x[0];
    d[0] = x[64];
    s[1] = x[128];
    s[mid] = x[128*mid];
    d[mid] = x[128*mid +64];

15
    d[0] = (d[0]<<1)-s[0]-s[1];

    s[0] = s[0]+(d[0]>>2);

    d_tmp0 = d[0];

20    s_tmp0 = s[1];

    for (ii=1; ii<mid; ii++) {
        d[ii] = (x[128*ii+64]<<1) - s_tmp0 - x[128*(ii+1)];
        s[ii] = s_tmp0 + ((d_tmp0+d_tmp1)>>3);
25    d_tmp0 = d[ii];
        s_tmp0 = s[ii];
    }

```



```

for (ii=1; ii<mid; ii++) {

    x[ii*64] = s[ii];
    x[(ii+mid+1)*64] = d[ii];
}
5  d[mid] = (d[mid]<<1) - (s[mid]<<1);
    s[mid] = s[mid] + ((d[mid-1]+d[mid])>>3);

    x[0] = s[0];
    x[(mid+1)*64] = d[0];
10  x[mid*64] = s[mid];
    x[(2*mid+1)*64] = d[mid];
}
}

```

- 15 After loop peeling, the only change ~~on~~with respect to the parameters is the vector length.
~~The~~Accordingly, the tables becomeare changed to the following:

Parameter	Value
Vector length	mid-2
Reused data set size	-
I/O IRAMs	6
ALU	2
BREG	0
FREG	2
Data flow graph width	2
Data flow graph height	6
Configuration cycles	6+(mid-2)

Parameter	Value
Vector length	mid-2
Reused data set size	-
I/O IRAMs	6
ALU	0
BREG	0
FREG	0

Data flow graph width	2
Data flow graph height	1
Configuration cycles	mid-2

The fusion of the inner loops is legal as there would be no loop-carried ~~dependences~~dependencies between the instructions formerly in the second loop and the instructions formerly in the first loop. ~~We obtain the~~The following loop nest may be
5 obtained.

```

for (nt=64 ; nt>= 16; nt >>=1) {
    for (i=0; i<nt*64 /*tmp_nt*/; i+=64) {

10      x = &int-_data[i];
        mid = (nt>>1)-1;

        s[0] = x[0];
        d[0] = x[1];
15      s[1] = x[2];
        s[mid] = x[2*mid];
        d[mid] = x[2*mid+1];

        d[0] = (d[0]<<1)-s[0]-s[1];
20      s[0] = s[0]+(d[0]>>2);

        d_tmp0 = d[0];
        s_tmp0 = s[1];

25      for (ii=1; ii<mid; ii++) {
            d[ii] = ((x[2*ii+1])<<1) - s_tmp0 - x[2*ii+2];
            s[ii] = s_tmp0 + ((d_tmp0+d[ii])>>3);
            d_tmp0 = d[ii];
            s_tmp0 = s[ii];
30      x[ii+mid+1] = d[ii];
    }

```

```

    d[mid] = (d[mid]-s[mid])<<1;
    s[mid] = s[mid]+((d[mid-1]+d[mid])>>3);

    x[0] = s[0];
5    x[mid+1] = d[0];
    x[mid] = s[mid];
    x[2*mid+1] = d[mid];

}
for (i=0; i<nt; i++) {
10
    x = &int_data[i];
    mid = (nt>>1)-1;

    s[0] = x[0];

15    d[0] = x[64];

    s[1] = x[128];

    s[mid] = x[128*mid];

    d[mid] = x[128*mid+64];

20    d[0] = (d[0]<<1)-s[0]-s[1];

    s[0] = s[0]+(d[0]>>2);

    d_tmp0 = d[0];

    s_tmp0 = s[1];

25

    for (ii=1; ii<mid; ii++) {

```

```

    d[ii] = (x[128*ii+64]<<1) - s_tmp0 - x[128*(ii+1)];
    s[ii] = s_tmp0 + ((d_tmp0 + d[ii])>>3);
    d_tmp0 = d[ii];
    s_tmp0 = s[ii];
5    x[ii*64] = s[ii];
    x[(ii+mid+1)*64] = d[ii];
}
d[mid] = (d[mid]<<1) - (s[mid]<<1);
s[mid] = s[mid] + ((d[mid-1]+d[mid])>>3);
10
x[0] = s[0];
x[(mid+1)*64] = d[0];
x[mid*64] = s[mid];
x[(2*mid+1)*64] = d[mid];
15 }
}

```

After loop fusion, ~~wethere are only have two loops, that~~ and they have the following same parameter table.

Parameter	Value
Vector length	mid-2
Reused data set size	-
I/O IRAMs	8
ALU	6
BREG	0
FREG	2
Data flow graph width	2
Data flow graph height	6
Configuration cycles	6+(mid-2)

When performing value range analysis, the compiler finds that *nt* ranges ~~take~~ take the values 64, ~~32, 32~~, and 16. The upper bound of the inner loops is *mid*, which depends on the value of *nt*. The analysis finds then that *mid* can take the values: ~~31, 15~~ 31, 15, and 7. Loops with constant loop bounds can be handled more efficiently on the PACT XPP. This means that the

inner loops can be better optimized if *mid* is replaced by a constant value. This will happen when the outer loop is unrolled. This way ~~we will obtain a bigger~~, a larger set of code will be obtained, but with 3 instances of the loop nest, each being a candidate for a configuration.

This can be seen as a kind of temporal partitioning. Thus, the outer loop is completely

5 unrolled giving six new loop nests.

```

for (i=0; i<4096; i+=64) { / *nt=64 */
    x = &int_data[i];
    mid=31;
10
    s[0] = x[0];
    d[0] = x[1];
    s[1] = x[2];
    s[31] = x[61];
15    d[31] = x[63];

    d[0] = (d [0] <<1) -s [0]-s [1] ;
    s[0] = s[0]+(d[0]>>2);

20    d_tmp0 = d[0];
    s_tmp0 = s[1];

    for (ii1; ii<31; ii++) {
        d[ii] = ((x [2*ii+1]) <<1) - s_tmp0 - x [2*ii+2];
25        s[ii] = s_tmp0+((d _tmp0 + d[ii] )>>3);
        d _tmp 0 = d [ ii ] ;
        s_tmp0 = s[ii];
        x[ii] = s[ii];
        x[ii+32] = d[ii];
30    }
    d[31] = (d [31] -s [31] ) <<1;
    s [31] = s [31] + ((d[30]+d [31] ) >>3) ;
    x[0] = s [0];

```

```

    x[32] = d[0];
    x[31] = s[31]
    x[63] = d[31];
}
5
for (i=0; i<64; i++) {

    x = &int_data[i];
    mid = 31;
10
    s[0] = x [0] ;
    d[0] = x [64];
    s[1] = x[128];
    s[31] = x[3968];
15
    d[31] = x[4032] ;

    d[0] = (d[0]<<1)-s[0]-s[1] ;
    s[0] = s[0]+(d[0]>>2) ;

20
    d_tmp0 = d[0];

    s_tmp0 = s[1];

    for (ii=1; ii<31; ii++) {
25
        d[ii] = (x[128*ii+64]<<1) - s_tmp0 - x[128* (ii+-1)];
        s[ii] = s_tmp0 + ((d_tmp0 + d[ii]) >>3) ;
        d_tmp0 = d[ii];
        s_tmp0 = s[ii];
        x[ii*64] = s[ii] ;
30
        x[(ii+32) *64] = d[ii];
    }
    d[31] = (d[31]<<1) - (s[31]<<1);
    s [31] = s[31] + ((d[30]+d[31] )>>3);

```

```

    x[0] = s [0];
    x[2048 ] = d[0] ;
    x[1984 ] = s[31];
5    x[4032 ] = d[31];
    }

    for (i=0; i<2048; i+=64) { /*nt = 32 */

10    x = &int _data [i] ;
        mid = 15 ;
        s[0] = x[0];
        d[0] = x[1];
        s [1] = x[2];
15    s[15] = x[30];
        d[15] = x [31] ;

        d[0] = (d [0]-<<1)-s[0]-s[1] ;
        s [0] = s [0] + (d[0]>>2) ;

20    d_tmp 0 = d [0];
        s_tmp0 = s[1];

        for (ii=1; ii<15; ii++) {
25    d[ii] = ((x[2*ii+1]) <<1) - s_tmp0 - x [2* ii+2] ;
        s [ii] = s_tmp0 + ((d_tmp0+d [ii]) >>3) ;
        d_tmp0 = d[ii];
        s_tmp0 = s[ii];
        x[ii]= = s[ii];
30

        x[ii+16] = d[ii];

        }
        d[15] = (d[15]-s[15])<<1;

```

```

s[15] = s[15]+((d[14]+d[15])>>3);

x[0] = s[0];
x[16] = d[0];
5  x[15] = s[15];
   x[31] = d[15];
}

for (i=0; i<32; i++) {
10
   x = &int_data[i];
   mid = 15;

   s [0] = x[0];
15  d[0] = x[64];
   s[1] = x[128];
   s[15] = x[1920] ;
   d[15] = x[1984];

20  d[0] = (d[0]<<1)-s[0]-s [1] ;
   s[0] = s [0] + (d [0]>>2) ;

   d_tmp0 = d[0];
   s_tmp0 = s[1];
25
   for (ii=1; ii<15; ii++) {
       d [ii] = (x [128*ii+64]<<1) - s_tmp0 - x[128*(ii+1)];
       s[ii] = s_tmp0 + ((d_tmp0+d[ii])>>3;
       d_tmp0 = d[ii];
30  s_tmp0 = s[ii];
       x[ii*64] = s[ii];
       x[ (ii+16)*64] = d[ii] ;
   }
   d[15] = (d[15]<<1)-(s[15] <<1) ;

```



```

s[15] = s[15]+( (d[14]+d [15] )>>3) ;

x[0] = s[0];
x[1024] = d[0];
5  x[960] = s [15] ;
    x[1984] = d[15];
}

for (i=0; i<1024; i+=64) { /* nt = 16*/
10
    x = &int_data [i] ;
    mid = 7 ;

    s[0] = x[0];
15  d[0] = x[1];
    s[1] = x[2];
    s[7 ] = x[14];
    d[7 ] = x[15] ;

20  d[0] = (d[0]<<1)-s[0]-s[1];

    s[0] = s[0]+(d[0]>>2);

    d_ tmp0 = d[0];
    s_ tmp0 = s[1];
25
    for (ii=1; ii<7; ii++) {
        d [ii] = ((x [2*ii+1]) <<1)- s_ tmp0 - x[2*ii+2] ;
        s[ii] = s_ tmp0+((d_ tmp0+d[ii])>>3);
        d_ tmp0 = d[ii];
30  s_ tmp0 = s[ii];
        x[ii] = s[ii];
        x[ii+8] = d[ii];

```

```

    }
    d[7] = (d[7]-s [7]) <<1;
    s[7] = s[7]+((d[6]+d[7])>>3);

5   x[0] = s [0] ;
    x[8] = d[0] ;
    x[7] = s[7];
    x[15] = d[7] ;
    }
10  for (i=0; i<16; i++) {

    x = &int _data [i];
    mid = 7;

15  s[0] = x[0];
    d[0] = x[64];
    s [1] = x[128];
    s[7] = x[896];
20  d[7] = x[960];

    d[0] = (d[0]<<1)-s[0]-s[1];
    s[0] = s[0]+(d[0]>>2);

25  d _tmp0 = d[0];
    s _tmp0 = s[1];

    for (ii=1; ii<7; ii++) {
        d[ii] = (x[128*ii+64]<<1) - s _tmp0 - x[128*(ii+1)];
30  s [ii] = s _tmp0 + ((d _tmp0+d[ii] )>>3)
        d _tmp0 = d[ii];
        s _tmp0 = s[ii];
        x[ii*64]= s[ii];
        x[(ii+8)*64] = d[ii];

```

```

}
d[7] = (d[7]<<1) - (s [7]<<1);
s [7] = s[7] + ((d[6]+d[7])>>3);

```

```

5  x[0] = s[0] ;
   x[512] = d[0];
   x [448] = s [7];
   x[960] = d[7];
}

```

10

In the parameter table, the vector length is the only value that ~~change~~. ~~We give it~~changes.

Below is a parameter table for the first two loops. To deduce the table for the other loops, the vector length has to be set to 14 and 66, respectively.

Parameter	Value
Vector length	30
Reused data set size	-
I/O IRAMs	8
ALU	6
BRED	0
FREG	2
Data flow graph width	2
Data flow graph height	6
Configuration cycles	6+30=36

15 5.6.3—Optimizing the Inner Loops

The efforts are then concentrated on the six inner loops. ~~In fact, if we look at them, they~~They all need 2 input data and 4 output 4-data. 2 more data are needed for the first iteration.

Hence ~~we need~~, at most, 8 IRAMs are required for the first iteration and 6 for the others.

This means that ~~we~~the loops can ~~unroll the loops~~be unrolled ~~twice needing~~, requiring 14

20 IRAMs for one iteration of the new loop bodies. Below ~~we present~~are presented only the unrolled inner loops ~~for commodity reasons~~.

FirstThe first loop may be as follows:

```

for (ii =1; ii<31; ii=ii+2) {
25   d[ii] = ((x[2* ii+1]) <<1) - s_tmp0 - x[2*ii+2];

```

```

s[ii] = s_tmp0 + ((d_tmp0+d [ii])>>3);

d_tmp0 = d[ii];
s_tmp0 = s[ii];
5  x[ii+1] = s[ii];
x[ii+33] = d[ii] ;
d[ii+1] = ((x[2* (ii+1)+1])<<1) - s_tmp0 - x[2*(ii+1)+2];
s [ii+1] = s_tmp0 + ((d_tmp0+d[ii+1]) >>3) ;
d_tmp0 = d[ii+1];
10 s_tmp0 = s [ii+1] ;
x[ii+1] = s[ii+1];
x[ii+33] = d[ii+1] ;
}

```

15 ~~Second~~The second loop may be as follows:

```

for (ii=1; ii<31;    ii=ii+2) {
    d[ii] = (x[128*ii+64]<<1) - s_ tmp0 - x[128*(ii+1)];

    s[ii] = s_tmp0 + ((d_tmp0+d[ii])>>3);
20

d_tmp0 = d[ii];

s_tmp0 = s[ii];
x[ii*64] = s[ii];
x[(ii+32) *64] = d[ii];
25 d[ii+1] = (x[128* (ii+1)+64] <<1) - s_tmp0 - x[128*
(ii+2)];
s[ii+1] = s_tmp0 + ((d_tmp0+d[ii+1])>>3);
d_tmp0 = d[ii+1];
s_tmp0 = s[ii+1];
30 x[ (ii+1) *64] = s [ii+1];
x[ (ii+33) *64] = d[ii+1];
}

```

~~Third~~The third loop may be as follows:

```
for (ii=1; ii<15 ; ii=ii+2) {
    d[ii] = ((x [2*ii+1]) <<1) - s _tmp 0 - x [2*ii+2] ;
5    s-[ii] = s_tmp0 + ((d_tmp0+d[ii])>>3);
    d_tmp0 = d[ii];
    s_tmp0 = s[ii];
    x[ii] = s[ii];
    x[ii+16] = d[ii] ;
10    d[ii+1] = ((x[2*(ii+1)+1])<<1) - s_tmp0 - x[2*(ii+1)+2];
    s [ii+1] = s_tmp0 + ((d_tmp0+d[ii+1])>>3);
    d_tmp0 = d[ii+1];
    s_tmp0 = s[ii+1];
    x[ii+1] = s[ii+1];
15    x[ii+17] = d [ii+1];
}
```

~~Fourth~~The fourth loop may be as follows:

```
for (ii=1; ii<15 ; ii=ii+2) {
20    d[ii] = (x[ 128*ii+64]<<1) - s_tmp0 - x[128*(ii+1)];
    s[ii] = s_tmp0 + ((d_tmp0 + d[ii]) >>3);
    d_tmp0 = d[ii];
    s_tmp0 = s[ii];
    x[ii*64] = s [ii];
25    x[(ii+16) *64] = d[ii];
    d[ii+1] = ( x[128*-(ii+1)+64]<<1) - s_tmp0 - x[128* (ii+2)];
    s [ ii ] = s_tmp0 + ((d_tmp0+d[ii+1]) >>3) ;
    d_tmp0 = d[ii+1];
    s_tmp0 = s[ii+1];
30    x[(ii+1) * 64 ] = s-[ii+1];
    x[(ii+17) *64 = d[ii+1] ;
}
```

~~Fifth~~ The fifth loop may be as follows:

```
for (ii= 1; ii<7; ii=ii+2) {
    d[ii] = ((x [2*ii+1]) <<1) - s_tmp0 - x [2*ii+2] ;
    s[ii] = s_tmp0 + ((d_tmp0+d[ii]) >>3) ;
5    d_tmp0 = d[ii]
    s_tmp0 = s[ii];
    x[ii] = s [ii];
    x[ii+8] = d[ii];
    d[ii+1] = ((x[2*(ii+1)+1])<<1) - s_tmp0 - x[2*(ii+1)+2];
10    s[ii+1] = s_tmp0 + ((d_tmp0+d[ii+1])>>3);
    d_tmp0 = d[ii+1];
    s_tmp0 = s[ii+1];
    x[ii+1] = s[ii+1];
    x[ii+9] = d[ii+1] ;
15 }
```

~~Sixth~~ The sixth loop may be as follows:

```
for (ii=1; ii<7; ii=ii+2) {
    d[ii] = (x[128*ii+64]<<1) - s_tmp0 - x[128*(ii+ 1)];
20    s[ii] = s_tmp0 + ((d _tmp0 + d[ii])>>3);
    d_tmp 0 = d[ii];
    s_tmp0 = s[ii] ;
    x [ii*64 ] = s [ii] ;
    x[ (ii+8) *64] = d[ii];
25    d[ii+1] = (x[128* (ii+1)+64 ]<<1) - s_tmp0 - x[128*(ii+2)];
    s [ii] = s_ tmp0 + ((d_-tmp0+d[ii+1]) >>3) ;
    d _tmp0 = d[ii+1];
    s_tmp0 = s[ii+1];
    x[(ii+1)*64] = s[ii+1];
30    x[(ii+9) *64] = d[ii+1];
}
```

~~We obtain the following~~ Fig. 53 is a dataflow graph of these loop bodies after a step of tree
balancing has been performed. We represent here only the graph corresponding The dataflow

graph of Fig. 53 corresponds to the first loop. To obtain the graphs for the other loops, only the input and output data need to be changed.

Each input and output data will occupy an IRAM. $d0$ and $s0$ will be the only values in their IRAM, enabling then the merge operations to select between $d0$, resp. 0 and $s0$ at the first iteration and the feedback values for the other iterations. Once the pipeline is filled, 8 values can be output in a cycle, corresponding to 4 values for array x . The same configuration is used for all loops; only the data in the IRAMs differ. ~~We give now~~ Below are result tables for only for the 2 first loops. The tables for the other tables loops are the same.

10

For the first two loops ~~we obtain~~, the following table is obtained, and the expected speedup with respect to a standard superscalar processor with 2 instructions issued per cycle is 1 5.3.

Parameter	Value
Vector length	30
Reused data set size	-
I/O IRAMs	14
ALU	12
BREG	0
FREG	2
Data flow graph width	2
Data flow graph height	10
Configuration cycles	10+15=25

Ops	Number
LD/ST (2 cycles)	14
ADDRCOMP (1 cycle)	2
ADD/SUB (1 cycle)	17
MUL (2 cycles)	0
SHIFT (1 cycle)	4
Cycles per iteration	51
Cycles needed for the loop (2-way)	$(51*15)/2=383$

15

References

~~The limitations of conventional processors are becoming more and more evident. The growing importance of stream-based applications makes coarse-grain dynamically~~

reconfigurable architectures an attractive alternative [3], [4], [6], [7]. They combine the performance of ASICs, which are very risky and expensive (development and mask costs), with the flexibility of traditional processors [5].

5 In spite of the possibilities we have today in VLSI development, the basic concepts of microprocessor architectures are the same as 20 years ago. The main processing unit of modern conventional microprocessors, the datapath, in its actual structure follows the same style guidelines as its predecessors. Although the development of pipelined architectures or supersealar concepts in combination with data and instruction caches increases the performance of a modern microprocessor and allows higher frequency rates, the main concept
10 of a static datapath remains. Therefore, each operation is a composition of basic instructions that the used processor owns. The benefit of the processor concept lays in the ability of executing strong control dominant application. Data or stream oriented applications are not well suited for this environment. The sequential instruction execution isn't the right target for that kind of applications and needs high band width because of permanent retransmitting
15 of instruction/data from and to memory. This handicap is often eased by using of caches in various stages. A sequential interconnection of filters, which do the according data manipulating without writing back the intermediate results would get the right optimisation and reduction of bandwidth. Practically, this kind of chain of filters should be constructed in a logical way and configured during runtime. Existing approach to extend instruction sets
20 use static modules, not modifiable during runtime.

Customized microprocessors or ASICs are optimized for one special application environment. It is nearly impossible to use the same microprocessor core for another application without losing the performance gain of this architecture.

25 A new approach of a flexible and high performance datapath concept is needed, which allows to reconfigure the functionality and make this core mainly application independent without losing the performance needed for stream based applications.

This contribution introduces a new concept of loosely coupled implementation of the dynamic reconfigurable XPP architecture from PACT Corp. into a static datapath of the SPARC compatible LEON processor. Thus, this approach is different from those, where the
30 XPP operates as a completely separate (master) component within one Configurable System-on-Chip (CsoC), together with a processor core, global/local memory topologies and efficient

multi-layer Amba bus interfaces [11]. Here, from the programmers point of view the extended and adapted datapath seems like a dynamic configurable instruction set. It can be customized for a specific application and accelerate the execution enormously. Therefore, the programmer has to create a number of configurations, which can be uploaded to the XPP-Array at run time, e. g. this configuration can be used like a filter to calculate stream-oriented data. It is also possible, to configure more than one function in the same time and use them simultaneously. This concept, promises an enormously performance boost and the needed flexibility and power reduction to perform a series of applications very effective.

6.2 — 1. LEON RISC Microprocessor

- 10 For implementation of this concept we chose the 32-bit SPARC V8 compatible microprocessor [1] [2], LEON. This microprocessor is a synthesisable, free available VHDL model which has a load/store architecture and has a five stages pipeline implementation with separated instruction and data caches.

Figure 66: LEON Architecture Overview

- 15 As shown in Figure 66 the LEON is provided with a full implementation of AMBA 2.0 AHB and APB on-chip bus, a hardware multiplier and divider, programmable 8/16/32-bit memory controller for external PROM, static RAM and SDRAM and several on-chip peripherals such as timers, UARTs, interrupt controller and a 16-bit I/O port. A simple power down mode is implemented as well.
- 20 LEON is developed by the European Space Agency (ESA) for future space missions. The performance of LEON is close to an ARM19-series but don't have a memory management unit (MMU) implementation, which limits the use to single memory space applications. In Figure 67 the datapath of the LEON integer unit is shown.

25 *Figure 67: LEON Pipelined Datapath Structure*

2. eXtreme Processing Platform — XPP

The XPP architecture [6], [7], [8] is based on a hierarchical array of coarse grain, adaptive computing elements called *Processing Array Elements (PAEs)* and a packet-oriented

~~communication network. The strength of the XPP technology originates from the combination of array processing with unique, powerful run-time reconfiguration mechanisms. Since configuration control is distributed over a Configuration Manager (CM) embedded in the array, PAEs can be configured rapidly in parallel while neighboring PAEs are processing~~
5 ~~data. Entire applications can be configured and run independently on different parts of the array. Reconfiguration is triggered externally or even by special event signals originating within the array, enabling self-reconfiguring designs. By utilizing protocols implemented in hardware, data and event packets are used to process, generate, decompose and merge streams of data.~~

10 ~~The XPP has some similarities with other coarse-grain reconfigurable architectures like the KressArray [3] or Raw Machines [4] which are specifically designed for stream-based applications. XPP's main distinguishing features are its automatic packet handling mechanisms and its sophisticated hierarchical configuration protocols for runtime and self-reconfiguration.~~

15 ~~6.2.1 2.1~~ **Array Structure**

~~A CM consists of a state machine and internal RAM for configuration caching. The PAC itself (see top right wand side of Figure 69) contains a configuration bus which connects the CM with PAEs and other configurable objects. Horizontal busses carry data and events. They can be segmented by configurable switch objects, and connected to PAEs and special~~
20 ~~I/O objects at the periphery of the device.~~

~~A PAE is a collection of PAE objects. The typical PAE shown in Figure 69 (bottom) contains a BREG object (back registers) and an FREG object (forward registers) which are used for vertical routing, as well as an ALU object which performs the actual computations. The ALU performs common fixed-point arithmetical and logical operations as well as several~~
25 ~~special three input opcodes like multiply-add, sort, and counters. Events generated by ALU objects depend on ALU results or exceptions, very similar to the state flags of a classical microprocessor. A counter, e.g., generates a special event only after it has terminated. The next section explains how these events are used. Another PAE object implemented in the XPP is a memory object which can be used in FIFO mode or as RAM for lookup tables,~~
30 ~~intermediate results etc. However, any PAE object functionality can be included in the XPP architecture.~~

~~6. 2. 2~~

~~6. 2. 3~~ — ~~2. 2~~ **Packet Handling and Synchronization**

PAE objects as defined above communicate via a packet-oriented network. Two types of packets are sent through the array: data packets and event packets. Data packets have a uniform bit width specific to the device type. In normal operation mode, PAE objects are self-synchronizing. An operation is performed as soon as all necessary data input packets are available. The results are forwarded as soon as they are available, provided the previous results have been consumed. Thus it is possible to map a signal flow graph directly to ALU objects. Event packets are one bit wide. They transmit state information which controls ALU execution and packet generation.

~~6.2. 4~~ — ~~2. 3~~ **Configuration**

Figure 69: Structure of an XPP device

Every PAE stores locally its current configuration state, i.e. if it is part of a configuration or not (states „configured” or „free”). Once a PAE is configured, it changes its state to „configured”. This prevents the CM from reconfiguring a PAE which is still used by another application. The CM caches the configuration data in its internal RAM until the required PAEs become available.

While loading a configuration, all PAEs start to compute their part of the application as soon as they are in state „configured”. Partially configured applications are able to process data without loss of packets. This concurrency of configuration and computation hides configuration latency.

~~6. 2. 5~~ — ~~2. 4~~ **XPP Application Mapping**

The Native Mapping Language (NML), a PACT proprietary structural language with reconfiguration primitives, was developed by PACT to map applications to the XPP array. It gives the programmer direct access to all hardware features.

In NML, configurations consist of modules which are specified as in a structural hardware description language, similar to, for instance, structural VHDL, PAE objects are explicitly allocated, optionally placed, and their connections specified. Hierarchical modules allow

component reuse, especially for repetitive lay outs. Additionally, NML includes statements to support configuration handling. A complete NML application program consists of one or more modules, a sequence of initially configured modules, differential changes, and statements which map event signals to configuration and prefetch requests. Thus configuration handling is an explicit part of the application program.

A complete XPP Development Suite (XDS) is available from PACT. For more details on XPP-based architectures and development tools see [6].

6. 3 — 3. LEON Instruction Datapath Extension

The system is designed to offer a maximum of performance. LEON and XPP should be able to communicate with each other in a simple and high performance manner, while the XPP is a dataflow orientated device, the LEON is a general purpose processor, suitable for handling control flow [1], [2]. Therefore, LEON is used for system control. To do this, the XPP is integrated into the datapath of the LEON integer unit, which is able to control the XPP.

Figure 71: Extended Datapath Overview

Due to unpredictable operation time of the XPP algorithm, integration of XPP into LEON data path is done in a loosely coupled way (Figure 71). Thus the XPP array can operate independent from the LEON, which is able to control and reconfigure the XPP during runtime. Since the configuration of XPP is handled by LEON, the CM of the XPP is unnecessary and can be left out of the XPP array. The configuration codes are stored in the LEON RAM. LEON transfers the needed configuration from its system RAM into the XPP and creates the needed algorithm on the array.

To enable a maximum of independence of XPP from LEON, all ports of the XPP input ports as well as output ports are buffered using dual clock FIFOs. Dual clocked FIFOs are implemented into the IO Ports between LEON and XPP. To transmit data to the extended XPP-based datapath the data are passed through an IO Port as shown in Figure 5. In addition to the FIFO the IO Ports contain logic to generate handshake signals and an interrupt request signal. The IO Port for receiving data from XPP is similar to Figure 5 except that the reversed direction of the data signals. This enables that XPP can work completely independent from LEON as long as there are input data available in the input port FIFOs and free space for result data in the output port FIFOs. There are a number of additionally

features implemented in the LEON pipeline to control the data transfer between LEON and XPP.

When LEON tries to write to an IO Port containing a full FIFO or read from an IO Port containing an empty FIFO a trap is generated. This trap can be handled through a trap handler. There is a further mechanism—pipeline holding—implemented, to allow LEON holding the pipeline and wait, for free FIFO space during XPP write access respectively wait for a valid FIFO value during XPP read access. When using pipeline holding the software developer has to avoid reading from an IO Port with empty FIFO while the XPP, respectively the XPP input IO Ports, contains no data to produce outputs. In this case a deadlock will occur and the complete system has to be reseted.

XPP can generate interrupts for the LEON when trying to read a value from an empty FIFO port or to write a value to a full FIFO port. The occurrence of interrupts indicates, that the XPP array cannot process the next step because it has either no input values or it cannot output the result value. The interrupts generated by the XPP are maskable.

The interface provides information about the FIFOs. LEON can react the number of valid values the FIFO contains.

Figure 72: LEON to XPP dual clock FIFO

The interface to the XPP appears to the LEON as a set of special registers. (Figure 6). These XPP registers can be categorized in communication registers and status registers.

For data exchange the XPP communication registers are used. Since XPP provides three different types of communication ports, there are also three types of communication registers, whereas every type is splitted into an input part and an output part:

The data for the process are accessed through XPP data registers. The number of data input and data output ports as well as the data bit width depends on the implemented XPP array.

XPP can generate and consume events. Events are one bit signals; The number of input events and output events depends on the implemented XPP array again.

Configuration of the XPP is done through the XPP configuration register. LEON reads the required configuration value from a file stored in his system RAM and writes it to the XPP configuration register.

There are a number of XPP status register implemented to control the behavior and get status information of the interface. Switching between the usage of trap handling and pipeline holding can be done in the hold register. A XPP clock register contains a clock frequency ratio between LEON and XPP. By writing this register LEON software can set the XPP clock relative to LEON clock. This allows to adapt the XPP clock frequency to the required XPP performance and consequently to influence the power consumption of the system. Writing zero to the XPP clock register turns off the XPP. At last there is a status register for every FIFO containing the number of valid values actually available in the FIFO.

This status registers provides a maximum of flexibility in communication between LEON and XPP and enables different communication modes:

—— If there is only one application running on the system. at the time, software may be enveloped in pipeline hold mode. Here LEON initiates data read or write from respectively to XPP. If there is no value to read respectively no value to write, LEON pipeline will be stopped until read or write is possible. This can be used to reduce power consumption of the LEON part.

In interrupt mode, XPP can influence the LEON program flow. Thus, the IO Ports generates an interrupt depending on the actual number of values available in the FIFO. The communication between LEON and XPP as done in interrupt service routines.

—— Polling mode is the classical way to access the XPP. Initiated by a timer event LEON reads all XPP ports containing data and writes all XPP ports containing free FIFO space. Between these phases LEON can compute other calculations.

It is anytime possible to switch between this strategies within one application.

The XPP is delivered containing a configuration manager to handle configuration and reconfiguration of the array. In this concept the configuration manager is dispensable because the configuration as well as any reconfiguration is controlled by the LEON through the XPP configuration register. All XPP configurations used for an application are stored in the LEON's system RAM.

6. 4 — 4. Tool and Compiler Integration

The LEON's SPARC V8 instruction set [1] was extended by a new subset of instructions to make the new XPP registers accessible through software. These instructions are based in the SPARC instruction format but they are not conform to the SPARC V8 standard.

Corresponding to the SPARC conventions of a load/store Architecture the instruction subset can be splitted in two general types. Load/store instructions can exchange data between the LEON memory and the XPP communication registers. The number of cycles per instruction are similar to the standard load/store instructions of the LEON. Read/write instructions are used for communications between LEON registers. Since the LEON register set is extended

	LEON alone	LEON with XPP in IRQ Mode	LEON with XPP in Poll Mode	LEON with XPP in Hold Mode
Configuration of XPP	=====	71.308 ns 17.827 cycles	84.364 ns 21.091 cycles	77.976 ns 19.494 cycles
2D iDCT (8x8)	14.672 ns 3.668 cycles	3.272 ns 818 cycles	3.872 ns 968 cycles	3.568 ns 892 cycles

Table 2 Performance on IDCT (8x8)

by the XPP registers the read/write instructions are extended also to access XPP registers. Status registers can only be accessed with read/write instructions. Execution of arithmetic instructions on XPP registers is not possible. Values have to be written to standard LEON registers before they can be target of arithmetic operations. —

The complete system can still operate any SPARC V8 compatible code. Doing this, the XPP is completely unused.

The LEON is provided with the LECCS cross compiler system [9] standing under the terms of LGPL. This system consists of modified versions of the binutils 2.11 and gcc 2.95.2. To make the new instruction subset available to software developers, the assembler of the binutils has been extended by a number of instructions according to the implemented instruction subset. The new instructions have the same mnemonic as the regular SPARC V8 load, store, read and write instructions. Only the new XPP registers have to be used as source respectively target operand. Since the modifications of LECCS are straightforward extensions, the cross compiler system is backward compatible to the original version. The

availability of the source code of LECCS has allowed to extend the tools by the new XPP operations in the described way.

The development of the XPP algorithms have to be done with separate tools, provided by PACT Corp.

5 6.5 — 5. — Application Results

As a first analysis application a inverse DCT applied to 8x8 pixel block was implemented. For all simulations we used 250 MHz clock frequency for LEON processor and 50 MHz clock frequency for XPP. The usage of XPP accelerates the computation of the IDCT about factor four, depending on the communication mode. However XPP has to be configured before computing the IDCT on it. Table 1 also shows the configuration time for this algorithm. As shown in Figure 74, the benefit brought by XPP rises with the number of IDCT blocks computed by it before reconfiguration, so the number of reconfigurations during complex algorithms should be minimised.

A first complex application implemented on the system is MPEG-4 decoding. The optimization of the algorithm partitioning on LEON and XPP is still under construction. In Figure 8 the blockdiagram of the MPEG-4 decoding algorithm is shown. Frames with 320 x 240 pixel was decoded. LEON by using SPARC V8 standard instructions decodes one frame in 23,46 seconds. In a first implementation of MPEG-4 using the XPP, only the IDCT is computed by XPP, the rest of the MPEG-4 decoding is still done with, LEON. Now, with the help of XPP, one frame is decoded in 17,98 s. This is a performance boost of more then twenty percent. Since the XPP performance gain by accelerating the idCT algorithm only is very low in the moment, we work on XPP implementations of Huffmann decoding, dequantisation and prediction decoding. So the performance boost of this concept against the standalone LEON will be increased.

25 6.6 — 6. — Conclusion

Today, the instruction datapaths of modern microprocessors reach their limits by using static instruction sets, driven by the traditional von Neumann or Harvard architectural principles. Away out of these limitations is a dynamic reconfigurable processor datapath extension achieved by integrating traditional static datapaths with the coarse grain dynamic reconfigurable XPP architecture (eXtreme Processing Platform). Therefore, a loosely

asynchronous coupling mechanism of the given instruction datapath has been developed and integrated onto a CMOS 0.13 μm standard cell technology from UMC. Here, the SPARC compatible LEON RISC processor is used, whereas its static pipelined instruction datapath has been extended to be configured and personalized for specific applications. This compiler compatible instruction set extension allows a various and efficient use, e.g. in streaming application domains like MPEG-4, digital filters, mobile communication modulation, etc. The introduced coupling technique by flexible dual clock FIFO interfaces allows asynchronous concurrency and adapting the frequency of the configured XPP datapath dependent on actual performance requirements, e.g. for avoiding unneeded cycles and reducing power consumption.

As represented above, the introduced concept combines the flexibility of a general purpose microprocessor with the performance and power consumption of coarse-grain reconfigurable datapath structures, nearly comparable to ASIC performance. Here, two programming and computing paradigms (control-driven von Neumann and transport-triggered XPP) are unified within one hybrid architecture with the option of two clock domains. The ability to reconfigure the transport-triggered XPP makes the system independent from standards or specific applications. This concept opens potential to develop multi-standard communication devices like software radios by using one extended processor architecture with adapted programming and compilation tools. Thus, new standards can be easily implemented through software updates. The system is scalable during design time through the scalable array-structure of the used XPP extension. This extends the range of suitable applications from products with less multimedia functions to complex high performance systems.

6. 7. — References

[1] — The SPARC Architecture Manual, Version 8, SPARC International INC.,
<http://www.spare.com>.

[2] — Jiri Gaisler: The LEON Processor User's Manual, <http://www.gaisler.com>

[3] — R. Hartenstein, R. Kress, and H. Reinig: A new FPGA architecture for word-oriented datapaths. In Proc. FPL'94, Prague, Czech Republic, September 1994. Springer LNCS 849

[4]—E. Waingold, M. Taylor, D. Srikarishna, V. Sarkar, W. Lee, V. Lee, J. Kim, M. Frank, and P. Finch. Bearing it all to software: Raw machines. *IEEE Computer*, pages 86–93, September 1997

[5]—J. Becker (invited Tutorial): Configurable Systems on Chip (CSoC); in: Proc. of 9th Proc. of XV Brazilian Symposium on Integrated Circuit, Design (SBCCI 2002), Porto Alegre, Brazil, September 5–9, 2002

[6]—PACT Corporation: <http://www.pactcorp.com>

[7]—The XPP Communication System, PACT Corporation, Technical Report 15, 2000

[8]—V. Baumgarte, F. Mayr, A. Nüchel, M. Vorbach, M. Weinhardt: PACT XPP—A Self-Reconfigurable Data Processing Architecture; The 1st Intl. Conference of Engineering of Reconfigurable Systems and Algorithms (ERSA ‘01), Las Vegas, NV, June 2001

[9]—LEON/ERC32 Cross Compilation System (LECCS), <http://www.gaisler.com/leecs.html>

[10]—W. Vorbach, J. Becker: Reconfigurable Processor Architectures for Mobile Phones; Reconfigurable Architectures Workshop (RAW 2003), Nice, France, April, 2003

[11]—J. Becker, M. Vorbach: Architecture, Memory and Interface Technology Integration of an Industrial/Academic Configurable System on Chip (CSoC); IEEE Computer Society Annual Workshop on VLSI (WVLSI 2003), Tampa, Florida, USA, February, 2003.

DATA PROCESSING

In embodiments of the present invention, support is provided for modern technologies of data processing and program execution, such as multi-tasking, multi-threading, hyper-threading, etc.

6—References

[1]—Markus Weinhardt and Wayne Luk. Memory Access Optimization for Reconfigurable Systems. *IEE Proceedings Computers and Digital Techniques*, 48(3), May 2001.

- [2] ~~Michael Wolfe. High Performance Compilers for Parallel Computing. Addison-Wesley, 1996.~~
- [3] ~~Hans Zima and Barbara Chapman. Superecompilers for parallel and vector computers. Addison-Wesley, 1991.~~
- 5 [4] ~~David F. Bacon, Susan L. Graham and Oliver J. Sharp. Compiler Transformations for High-Performance Computing. *ACM Computing Surveys*, 26(4):325-420, 1994.~~
- [5] ~~Steven Muchnick. Advanced Compiler Design and Implementation. Morgan Kaufmann, 1997.~~
- 10 [6] ~~João M. P. Cardoso and Markus Weinhardt. XPP-VC: A C Compiler with Temporal Partitioning for the PACT XPP Architecture. In *Proceedings of the 12th International Conference on Field-Programmable Logic and Applications, FPL'2002*, volume 2438 of LNCS, pages 864-874, Montpellier, France, 2002.~~
- [7] ~~Markus Weinhardt and Wayne Luk. Pipeline Vectorization. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 20(2):234-248, February 2001.~~
- 15 [8] ~~V. Baumgarte, F. May, A. Nüchel, M. Vorbach and M. Weinhardt. PACT XPP—A Self-Reconfigurable Data Processing Architecture. In *Proceedings of the 1st International Conference on Engineering of Reconfigurable Systems and Algorithms (ERSA'2001)*, Las Vegas, Nevada, 2001.~~
- 20 [9] ~~Katherine Compton and Scott Hauck. Reconfigurable Computing: A Survey of Systems and Software. *ACM Computing Surveys*, 34(2):171-210, 2002.~~
- [10] ~~Sam Larsen, Emmett Witchel and Saman Amarasinghe. Increasing and Detecting Memory Address Congruence. In *Proceedings of the 2002 IEEE International Conference on Parallel Architectures and Compilation Techniques (PACT'02)*, pages 1-8-29, Charlottesville, Virginia, September, 2002.~~
- 25 [11] ~~Daniela Genius and Sylvain Lelait. A Case for Array Merging in Memory Hierarchies. In *Proceedings of the 9th International Workshop on Compilers for Parallel Computers, CPC'01, Edinburgh, Scotland, June 27th-29th 2001.*~~

Second Major part

Another aspect of the present invention will be described in the following. This aspect deals with problems relating to the implementation of hyper-threading, multi-threading, multi-tasking, scheduling and execution of parts of configurations and so forth.

- 5 It is noted that WO 00/49496 already discloses a method for execution of a computer program, using a processor comprising a configurational unit capable of executing reconfigurable instructions in fact of which can be redefined at runtime.

A problem in conventional processor architectures exists if a coupling of for example sequential processors is needed and/or technologies such as a data streaming, hyper-
10 threading, multi-threading and so forth shall be used in a useful way enhancing performance. Techniques known in prior art, such as WO 02/50665 A1, do not allow for a sufficiently efficient way of effecting data exchange between the ALU of a CPU and the configurable data processing logic cell field, be it an FPGA, DSP or the like, that data exchange being effected via registers in the prior art. In other words, it is necessary to first sequentially write
15 data into a register and then retrieve them sequentially and restore them sequentially as well.

Another problem exists, if an external access to data is requested in known devices such as those cited used inter alia to implement functions in the configurable data processing logic cell field, DSP, FPGA or the like, that can not be processed sufficiently on the CPU-integrated ALU. Accordingly, the data processing logic cell field is practically used to allow
20 for user-defined opcodes that can process data more efficiently than is possible on the ALU of the CPU without further support by the data processing logic cell field. In the prior art, the coupling is generally word-based, not block-based. It is a further important aspect of the present invention that it has been realized that for data streaming data processing block-based coupling is highly preferable. At any rate, a more efficient data processing, in
25 particular more efficient than possible with a close coupling via registers, is highly preferable.

Another method for the use of logic cell fields consisting of coarse and/or fine granular logic cells and logic cell elements consists in a very loose coupling of such a field to a conventional CPU and/or a CPU core in embedded systems. Here, a conventional sequential
30 program can be executed on the CPU, for example a program written in C, C++ or the like, wherein the instantiation of the data stream processing by the fine and/or coarse granular

data processing logic cell field is effected via that sequential program. Then, the problem exists that in programming said logic cell field a program not written in C or any other sequential high-level language must be provided for the data stream processing. It would be preferable here to allow for C programs to run both on a conventional CPU architecture as well as on the data processing logic cell field operated therewith; in particular despite the fact that a quasi-sequential program execution should maintain the capability of data streaming in the data processing logic cell fields whereas simultaneously the capability exists to operate the CPU in a not too loosely coupled way.

It is already known to provide for sequential data processing within a data processing logic cell field, see for example DE 196 51 075, WO 98/26356, DE 196 54 846, WO 98/29952, DE 197 04 728, WO 98/35299, DE 1 99 26 538; WO 00/77652, DE 102 12 621. Here, partial execution is achieved within a single configuration, for example to reduce the amount of resources needed, optimize the time of execution and so forth; however, this does not lead automatically to allowing a programmer to translate or transfer high-level language code automatically onto a data processing logic cell field as is the case in common machine models for sequential processes. The compilation, transfer or translation of a high-level language code onto data processing logic cell fields according to the methods known for models of sequentially executing machines does remain difficult in fact.

In the prior art, it is further known that configurations that effect different functions on parts of the area respectively can be simultaneously executed on the processing array and that a change of one or some of the configuration(s) without disturbing other configurations is possible at run time.

Methods and hardware implemented means for the implementation are known to ensure that the execution of partial configurations to be loaded onto the array is possible without deadlock. Reference is being made to DE 196 54 593, WO 98/31102, DE 198 07 872, WO 99/44147, DE 199 26538, WO 00/77652, DE 100 28 397, WO 02/13000. This technology allows in a certain way a certain parallelism and, given certain forms and interrelations of that configurations or partial configurations for a certain way of multitasking/multi-threading, in particular in such a way that the planning, that is the scheduling and/or the planning control for time use, can be provided for. Furthermore, from the prior art, time-use planning control means and methods are known per se, that, at least under a corresponding interrelation of configurations and/or assignment of configurations to certain tasks and/or

~~threads to configurations and/or sequences of configurations allow for a multi-tasking and/or multi-threading. The use of such time use planning control means used in the prior art for configuring and management of configurations for the purpose of scheduling of tasks, threads, multi and hyper threads is considered to be inventive per se.~~

- 5 ~~In at least a partial aspect of preferred variants, it is preferable to provide for support of modern technologies of data processing and program execution such as multi-tasking, multi-threading, hyper-threading and so forth.~~

Another important aspect~~In embodiments~~ of the present invention ~~can be seen in that,~~ data are inputted into ~~that~~the data processing logic cell fields in response to the execution of a load configuration by ~~that~~the data processing logic cell fields, and/or ~~in that~~ data are stored from ~~that~~the data processing logic cell fields by executing a store-configuration. Accordingly, it is preferred to provide the load- and/or store-configurations in such a way that the addresses of those memory cells used are directly or indirectly generated within the data processing logic cell fields, the addresses indicating those memory cells and/or locations to which an access
10 ~~has to be effected as a~~ load- and/or store-access, that is as i.e., a read- and/or write-access. By configuring address generators within the configuration it becomes possible to load a plurality of data into the data processing logic cell fields where they can be stored in ~~internal~~ RAMs (IRAMs) and/or within the internal cells such as EALUs having registers and/or in other dedicated memory and/or storage ~~means~~. The load- or store-configuration,
15 ~~respectively,~~ thus allows for a ~~block-wise~~blockwise and thus almost data-stream-like loading and storing of data, this being in particular much faster than a single access and can be executed prior to or during the execution of one or more data processing - and/or data handling in a data altering manner - configurations processing ~~that~~the preloaded data.

The data loading can take place, provided that that logic cell fields are, as is typically the
25 case, sufficiently large, - in small partial areas thereof, while other partial areas are executing other tasks. For example, in ~~the ping-pong-like data processing described in other published documents by the present applicant, that known~~PACT is discussed a ping-pong-like data processing ~~relying that relies~~ on memory cells provided on each side of ~~that~~the data processing field, ~~where data in.~~ In a first data processing step, data stream from the memory
30 on one side through the data processing field to the memory on the other side of ~~that~~the data processing field, ~~where they.~~ The data are stored there as intermediate results while, if necessary, the array is reconfigured, ~~the.~~ The intermediate results ~~are then streaming backstream~~ for further processing ~~and so forth, etc.~~ Here, a memory strip on one side and/or memory part on one side can be preloaded with data by a load configuration in one array part,
35 while ~~on~~in the memory part on the other side of the logic cell field data are written out using

a store-configuration. ~~It is to be noted that such~~ Such a simultaneous load-/store-way of data processing is possible even without spatial distribution and/or separation of memory areas in which data are retrieved and/or in which data are stored.

It is possible to effect the data loading from a cache and/or into a cache. ~~It is advantageous;~~
5 ~~if~~ In one embodiment, the external communication to large memory banks ~~is~~ may be handled via a cache controlling unit without having to provide for separate circuitry within the data-processing logic cell field ~~in that the~~. The access in a writing or reading manner to cache-memory-means typically is very fast and has a small latency (if any). ~~Also, and in that also~~ typically a CPU-Unit is, ~~here typically~~ for example, via a load-/store-unit, coupled to the
10 cache so that access to data and an ex-change thereof between the CPU-core and the data processing logic cell fields can be effected ~~fast~~ quickly, block-wise, and such that not every single datum needs to be transferred via a separate instruction that must be fetched, for example, by the opcode-fetcher of the CPU and processed therein.

This cache-coupling ~~is also highly preferred compared to~~ may be much better than the
15 coupling of the data processing logic cell field to the ALU ~~With~~ with the CPU via registers, if those registers communicate only via a load-/store-unit with the cache, as is ~~known in the non-PACT prior art.~~ the conventional case.

~~It is possible to provide for~~ In an embodiment of the present invention, a further data connection may be provided to and/or from the load-/store-unit of the, or one of ~~that the~~,
20 sequential-CPU-units ~~connecting~~ connected to the data processing logic cell fields and/or - their registers.

~~It is to be mentioned that it is possible to address~~ possible to address units via separate input-/output ports of the data processing logic cell field, which can in particular be provided as a VPU or XPP, and/or to ~~address that~~ address the data processing logic cells via one or more
25 multiplexers downstream a single port.

~~It is also to be mentioned that beside of the block-wise~~ Besides the blockwise and/or streaming and/or random mode access to cache areas in a writing and a reading manner and/or to the load-/store-unit ~~anchor~~ and/or the ~~per se~~ known connection to the registers of a sequential CPU ~~it is also possible to provide for a connection,~~ in an embodiment of the
30 present invention, a connection is provided to an external mass memory such as a RAM, a ~~harddisc~~ hard disc or any other data exchange or input or output port such as an antenna ~~and so forth.~~ It is possible to provide for, etc. In an embodiment, separate ports may be provided for the access to several of such units and/or memory means. ~~It is also to be mentioned that~~ suitable Suitable drivers, signal conditioning circuitry, and so forth; ~~should~~ may accordingly
35 be provided. Furthermore, ~~it is to be mentioned that in particular~~, although not exclusively

for the handling of a data stream streaming into ~~that~~the data processing logic cell field and/or out of ~~that~~the data processing logic cell fields, the logic cells of ~~that~~the field can ~~comprise~~include ALUs or EALUs, respectively, which can have at their input- and/or output ports short, fine-granularly configurably FPGA-like circuitries, for example, to cut out 4-bit-

5 blocks out of a continuous data stream as is necessary, for example, for an MPEG-4-decoding. This ~~is~~may be advantageous ~~on one hand, for example,~~ if a data stream is to be ~~inputted~~input into the cell and is to be processed or preprocessed without blocking larger PAE-units. It ~~is in particular of advantage, if the ALU is~~In an embodiment of the present invention, the ALU may be provided- as an SIMD-ALU; ~~here, for.~~ For example, a very

10 broad data word having, for example, a broad 32-bit-data-width is may accordingly be split via an FPGA-like stripe in front of the SIMD-ALU ~~splitting the broad e.g., 32-bit-data into~~ eight data words having, for example, a 4-bit-data-width that can then be processed parallelly in the SIMD-ALU, increasing the overall performance of the system significantly, provided that the respect of applications are needed.

15 Furthermore, it is noted that when reference is being ~~had~~made to- FPGA-like pre- or post ~~that~~ structures, it is not absolute necessary to refer to 1-bit-granular devices; ~~in contrast.~~ Instead, it would be possible to provide finer-granular structures of a, for example, 4-bit-, instead of the hyper-fine-granular 1-bit, structures. In other words, the FPGA-like input- and/or output-structures, in front of or data downstream of ~~that~~the ALU-unit, ~~in.~~ In particular, SIMD-ALU-

20 units ~~are~~may be configurable in such a way that ~~always~~ 4-bit-data-words are always processed. It is also possible to provide for a cascading, so that, for example, incoming 32-bit-data-width words are separated into 4-bit parts by 8-bit FPGA-like structures in sequence of each other, then ~~processing that~~the four 8-bit data words are processed in four FPGA-like 8-bit-width structures, then ~~providing for~~ a second stripe of 8 separate 4-bit-wide FPGA-like

25 structures are provided, and, if necessary, ~~provide for example~~ sixteen separate having parallel 2-bit FPGA-like structures, for example, are provide. If this is the case, a significant reduction of the overhead compared to a hyper-fine-granular 1-bit FPGA-like structure can be achieved. This ~~allows to~~may allow for significantly ~~reduction~~reducing the configuration memory ~~and so forth, etc.,~~ thus saving on silicon area.

30 It is ~~to be noted that all~~many of thatthe coupling advantages ~~are achievable~~may be achieved using data block streams via a cache; ~~however.~~ However, it is preferred in particular if ~~that~~the cache is built slice-wise and if an access onto several slices ~~can take place~~ simultaneously, and in particular onto all slices ~~at the same time, can take place~~ simultaneously. It ~~is~~may be advantageous if, ~~as will be understood,~~ the data processing logic

35 cell field (XPP) and/or the sequential CPU and/or CPUs ~~have to process~~ a plurality of threads, ~~be it whether~~ by way of hyper-threading, multi-tasking, and/or multi-threading. It ~~is~~may also be preferable to provide cache-storage ~~means-with~~ slice access and/or slice access

enabling control means. For example, every single thread can be assigned a separate slice ~~thus, thereby~~ allowing that on processing that thread the respective cache areas are accessed on the re-entry of the group of codes to be processed. ~~However, it is to be understood that the cache needs~~ However, the cache need not necessarily be separated into slices and ~~that, every,~~ even if the cache is separated into slices, not every single thread must be assigned a separate slice. ~~However, it is to be noted that, although~~ this may be ~~[[is]]~~ a highly preferred method. Furthermore, it is to be noted that there may be cases ~~wherein~~ where not all cache areas are used simultaneously or temporarily at a given time. Instead, it is to be expected that in typical data processing applications, such as in hand-held mobile telephones, laptops, cameras and so forth, etc., there ~~will exist~~ may be periods during which not the entire cache is needed. ~~Thus~~ Accordingly, it ~~is~~ may be highly preferred ~~if advantageous that~~ certain cache-areas can be separated from the power source in such a way that² the energy consumption is significantly reduced, in ~~particularly~~ particular, close to or exactly to 0. This can be achieved by a power supply separation means arrangement adapted to separate cache slices from power. The separation can either be effected by a down-clocking, separation of clock-lines, and/or the overall separation of a power supply. ~~It is in~~ In particular, it may be possible to provide for such a separation for every single cache slice, for example, by an access identification arrangement adapted to identify whether or not a thread, hyper-thread, task, or the like is currently assigned to a respective cache slice. In case the access identification means arrangement indicates and/or detects that this is not the case, ~~typically there may be a~~ separation of slice from a clock-line and/or even the power-line ~~can be envisaged as possible.~~ It is also ~~to be~~ noted that on repowering-up after a separation from power, it is possible to immediately access the cache area, ~~thus.~~ Thus, no significant delay by switching an ON or ~~[[OF]]~~ OFF of the power is to be expected, as long as the hardware is implemented with current semiconductor technologies.

~~A further particular advantage~~ In embodiments of the ~~present~~ present invention ~~resides in the fact that,~~ although the transfer of data and/or operands is possible in a block-wise manner, no particular balancing is needed to ensure that exactly the same times of execution, of data processing steps in the sequential CPU and the XPP and/or other data processing logic cell fields are achieved. Instead, the processing ~~is~~ may frequently be independent, in particular in such a way that the sequential CPU and the data processing logic cell field can be considered as separate resources by a scheduler. This allows for the immediate implementation of known data processing programs splitting technologies such as multi-tasking, multi-threading, and/or hyper-threading. ~~The advantage that~~ A result of a data path balancing is not being necessary has as a result is that, for example, in a sequential CPU a number of-pipeline stages may be included, clock frequencies and/or schemes of clocking may be achieved in a different way ~~and so forth, etc.~~ It is a particular advantage if asynchronous logic is needed. ~~A further advantage~~ In an embodiment of the present invention ~~results in that,~~ by configuring

a load- and a store-configuration into the data processing logic cell fields, the data inside the field can be loaded into that field or out of that field which is ~~no more~~not controlled by the clock frequency of the CPU, the performance of the opcode fetcher ~~and so forth, etc.~~ In other words, the opcode fetcher ~~is no more~~does not bottle-neck the data throughput to the data logic cell field without having an only loose coupling.

In a ~~particularly preferred~~an example embodiment of the present invention, it is possible to use the known CT or CM (commonly employed in the XPP-unit, also given the fact that with one or more, even hierarchically arranged XPP-fields having in ~~preferred~~some embodiments their own CTs while simultaneously using one or more sequential CPUs) ~~here~~ as a quasi hyper-threading hardware-management unit ~~having, which may have the inherent~~ advantage that known technologies, such as FILMO and ~~so forth~~others, become applicable for the hardware support and management of hyper-threading ~~and so forth; it, etc.~~ It is alternatively ~~and/or possible~~, in particular in a ~~hierarchically~~hierarchical arrangement ~~also possible~~, to provide the configurations from the opcode-fetcher of a sequential CPU via the coprocessing interface, allowing ~~to instantiate~~for instantiation of an XPP and/or data processing logic cell field call by the sequential CPU to effect data processing on the data processing logic cell field. Cache coupling and/or load and/or store configurations providing adress generators for loading and/or storing of data into ~~that~~the data processing logic cell field or out of that field ~~will~~may provide for the data exchange of the XPP. In other words, the coprocessor-like coupling to the data processing logic cell field ~~is~~may be enabled while, simultaneously, a data stream-like data loading is effected via cache- and/or I/O-port coupling.

~~It is to be mentioned that the~~The method of coprocessor coupling, that is the indicated coupling of the data processing logic cell field ~~will, may~~ typically result in the scheduling of ~~that~~the logic cell field taking place on the sequential CPU and/or a supervising scheduler unit and/or a respective scheduler means. In such a case, the threading control and/or ~~manage~~management practically takes place on the scheduler and/or the sequential CPU. Although this is ~~per se~~ possible, this will not ~~be necessarily~~ be the case where the ~~most~~ easiest implementation of the invention is sought. The data processing logic cell field can be called in a conventional manner, such as has been the case in a standard coprocessor such as a combination of ~~8086/8087~~. 8086/8087.

~~It is to be mentioned that in the particularly preferred~~In one example embodiment, independent of ~~the way of~~ its configuration, ~~e.g., as be~~ this a coprocessor interface, the configuration manager acting as scheduler at the same time or in any other way, it is possible to ~~adress~~address memory within or in an immediate vicinity of the data processing logic cell fields or under its management, in particular memory within the XPP-architecture ~~as is known by the applicant, RAM-PAEs or other; accordingly, etc.~~ Accordingly, managing

internal memories such as a vector register ~~is suggested, that is~~ may be advantageous. That is, the data volumes loaded via the load configuration ~~are~~ may be stored vector-like in vector registers in the internal-memory-cells, and thereafter accessing said registers may be accessed after loading and/or activating of a new configuration for effecting the actual data processing.

5 (It is ~~to be~~ noted that a data processing configuration can be referred to as one configuration even in a case that where several distinct configurations are to be processed simultaneously, one after the other or in a wave-like modus.)

~~Here, a~~ A vector register can be used to store results and/or intermediate results in the internal or internally managed memory cell elements. The vector register-like accessed memory
10 ~~means~~ in the XPP can be used also, after reconfiguration of the processing configuration by loading a store configuration in a suitable manner, in a way that takes place again in a data-stream-like manner, be it via an I/O-port directly streaming data into external memory areas and/or, ~~as particularly preferred,~~ into cache areas or out of these which then can be accessed at a later stage by the sequential CPU and/or other configurations- executed on the other data
15 processing logic cell field, ~~in-particularly in~~ a data processing logic cell field having produced said data in the first place.

In ~~a particularly preferred~~ one example embodiment, at least for certain data processing results and/or intermediate ~~re-suits; results,~~ for the memory and/or memory register ~~means~~ registers into which ~~that data~~ the processed data are to be stored, not an internal
20 memory, but instead a cache area having access reservation, ~~in-particularly~~ cache areas which are organized in a slice-wise manner, can be used. This can have the disadvantage of a larger latency, in particular if the paths between the XPP and/or data processing logic cell fields to or from the cache are of considerable length such that signal transmission delays need to be considered. Still, this ~~allows to avoid~~ may allow for additional store configurations to be
25 avoided. It is also ~~to be~~ noted, that ~~such a~~ this way of storing data in a cache area ~~be-~~ comes becomes, on the one hand, possible by placing the memory, into which data are stored, physically close to the cache controller and embodying that memory as a cache, but that alternatively and/or additionally the possibility exists to submit a part of a data processing logic cell field memory area, or internal memory, such as e.g., in the "RAM over PAE" case
30 under the control of one or several cache-memory controller(s)-, e.g., in the "RAM over PAE" case.

This ~~is~~ may be advantageous if the latency in storing the data processing results are to be kept small ~~in storing data processing results,~~ while latency in accessing the memory area serving as a quasi-cache to other units will not be too significant in other cases.

35 ~~It is also to be mentioned that~~ In an embodiment is possible in such a way that of the present invention, the cache controller of the known sequential CPU ~~addresses~~ may address as a cache

a memory area that is, without serving for the purpose of data exchange with a data processing logic cell field ~~is~~, physically placed onto that data processing logic cell field and/or close to that field. This ~~is~~ may be advantageous in that, if applications are run onto the data processing logic cell fields having a very small local memory need and/or if only few other configurations compared to the overall amount of memory space available are needed, these memory areas can be assigned to one or more sequential CPUs as cache or additional cache. ~~It is to be mentioned that if in~~ In such a case the cache controller may be adapted for the management of a cache area having a dynamically varying size.

A ~~dynamic~~ dynamic cache-size management and/or ~~dynamic~~ dynamic cache management size means for the ~~dynamic~~ dynamic cache management may ~~particularly~~ take into account the work load on the sequential CPU and/or the data processing logic cell fields. In other words, ~~it could be analyzed for example so as to enable fast reconfiguration (whether by way of wave-reconfiguration or in any other way),~~ how many NOPs in a given time unit are executed on the sequential CPU and/or how many configurations are preloaded in— the dynamically reconfigurable field in the memory areas provided therefore, ~~so as to enable fast reconfiguration (be it by the way of wave reconfiguration or in any other way).~~ The ~~dynamic~~ may be analyzed. The dynamic cache size or cache size management disclosed ~~herewith is, in a highly preferred manner, herein may be~~ runtime dynamical, ~~that~~. That is, the cache controller ~~controls~~ may control a momentary cache size ~~respectively~~ that can be changed from clock-cycle to clock-cycle or from one group of clock-cycles to ~~the~~ another. It is also ~~to be~~ noted that the access management of a data processing logic cell field with access as internal memory, such as vector register, is possible. While, as ~~previously~~ discussed above, a configuration management unit can be provided, it is ~~now to be~~ explicitly noted that such units and their way of operation, allowing in particular the preloading or configurations ~~currently~~ not yet needed, can be used very easily to effect the multi-task operation and/or hyper-threading and/or multi-threading, in particular for task- and/or thread- and/or hyper-thread switches. ~~Here, it is also noted that during~~ During the runtime of a thread or a task, it is possible to preload configurations for different tasks and/or threads and/or hyper-threads into the PAE-array. This ~~then may~~ allows to for a preload of configurations for a different task ~~and/or~~ thread if the current thread or task ~~can not~~ cannot be executed, for example because data ~~must be waited for, be it that~~ are awaited, whether where they have not yet been received, for example due to latencies, ~~be it because or where~~ a resource is blocked by another access. In case of the configuration preloading for a different task or thread, a switch or change becomes possible without the disadvantage of a timing overhead ~~because~~ due to the, for example, shadow-like loaded configuration execution.

It is in ~~principal~~ principle possible to use this technique also in cases where the most likely continuation of an execution is predicted and a prediction is missed. However, this way of

operation ~~will be most preferred~~ may be particularly advantageous in cases free of predictions. When using a pure sequential CPU and/or several pure sequential CPUs, the configuration manager thus also acts as and realizes a hyper-threading management hardware. ~~Reference is being had to DE 198 07 872, WO 99/44147 and WO 99/44120.~~ It can be considered as

5 sufficient, in particular in case where the CPU and/or several sequential CPUs ~~shall~~ have a hyper-threading management, to keep partial circuitry elements such as the FILMO; ~~described in the documents referred to, discussed in DE 198 07 872, WO 99/44147, and WO 99/44120.~~ In particular, the use of in an embodiment of the present invention, the configuration manager ~~described therein~~ discussed in these documents with and/or without
10 FILMO may be provided for use with the hyper-threading management for one and/or more purely sequential CPUs with; or without coupling to a data processing logic cell field ~~is disclosed herewith and claimed as inventive; although it is also to be noted that not in all cases where new and inventive features are disclosed in the present application these features are explicitly referred to as being inventive.~~

15 It ~~is also to be noted that~~ the plurality of CPUs can be realized with known techniques ~~such as known, for example from, such as those discussed~~ DE 102 12 624, 621 and PCT/EP 02/10572. It is also ~~to be noted that~~ DE 106 51 075, DE 106 54 846, DE 107 04 728; ~~728,~~ WO 98/26356, WO 98/~~29952~~ 29952, and WO 98/35299 ~~disclosed~~ discuss how to implement sequencers having ring- and/or random-access memory means in data processing logic cell
20 fields.

It is ~~to be noted that~~ a task-, thread- and/or hyper-thread switch ~~respectively~~ can be effected with the known CT-technology such that performance-slices and/or time-slices are assigned to a software implemented operating system scheduler by the CT, during which slices it is determined which parts of tasks and/or threads are subsequently to be executed provided that
25 resources are available.

~~An~~ The following is an example is given as follows:- First, an ~~address~~ address sequence is generated for a first task during which the execution of a load configuration loads data from a cache memory coupled ~~with~~ to the data processing logic cell field in the described manner. As soon as the data are present, the execution of a second configuration, the actual data-
30 processing configuration, can be started. This configuration can be preloaded as well; since it is certain that this configuration is to be executed provided that no interrupts or the like cause task switches. In conventional processes there is ~~now~~ the known ~~the~~ problem of the so -called cache-miss, ~~wherein~~ where data are requested that are not yet available in the cache. If such a case occurs in the coupling according to embodiments of the present invention, it is possible
35 to switch over to another thread, hyper-thread and/or task, in particular that has ~~in particular~~ been previously determined as the one to be executed next ~~by the~~, in particular by the

software implemented operating systems scheduler and/or ~~another~~other hard- and/or software implemented unit operating accordingly, and that has thus been preloaded in an available configuration memory of the data processing logic cell field, in particular preloaded in the background during the execution of another configuration, for example the load configuration which has effected the loading of data that are now ~~waited for~~ awaited.

It ~~to be~~ noted that it is possible to provide for separate configuration lines, these being, e.g., separate from communication lines used in the connection of ~~the~~, in particular, the coarse-granular data processing logic cells² of the data processing logic cell field. Then, if the configuration to which, due to the task-, thread-, and/or hyper-thread- switch, processing has been switched over, has been executed, and in particular has been in the preferable non-dividable, uninterruptable, and hence quasi atomic configuration executed until its end, a further other configuration as predetermined by that scheduler, ~~in-particularly~~ said operating system-like scheduler ~~is executed~~, and/or a configuration, for which the assigned load configuration has been executed may be executed. Prior to the execution of a processing configuration, for which a load configuration has been executed previously, a test can be ~~made~~performed to determine whether or not the respective ~~dated~~data have been streamed into the array, e.g., checking if the latency time which typically occurs has lapsed and/or the data are actually present.

In other words, latency times which occur as configurations are not yet preloaded, data have not yet been loaded, and/or data have not yet been stored, are bridged and/or covered by executing threads, hyper-threads, and/or tasks which have been preconfigured and which process data that are already available or can be written to resources that are available for writing thereto. In this way, latency times are covered and/or bridged and, provided a sufficient number of threads, hyper-threads, and/or tasks are to be executed, the data processing logic cell field ~~will~~can have an almost 100 % load.

In embodiments of the the present system invention, it is possible to realize a real time system despite the coupling of the array to a sequential CPU ~~and~~, in particular, while still having a data stream capability. In order to ensure real time capabilities it must be guaranteed that incoming data or interrupts ~~respectively signaling~~signaling incoming data are reacted upon without exceeding an allowed maximum time. This can be effected by causing a task switch on an interrupt and/or, for example, if the interrupts have a certain priority, by determining that a certain interrupt is currently to be ignored, which has to be determined within a certain time as well. A task switch in such systems capable of real time processing will thus typically be possible in one of three ~~ways namely either~~instances, which are when a task has run for a certain time (watch dog-principle), ~~not at non-~~at non-availability of a resource, ~~be it whether~~

due to a blockade, due to another access, or due to latencies, and/or ~~in the case of~~ the occurrence of interrupts.

A way of implementing one of these variants may ensure the real time capability. In a first alternative, one resource which is under the control of the CT ~~oder~~ scheduler switches over to processing the interrupt. If the allowed response time to a certain interrupt is so long that the configuration currently configured can- be executed without interruption this is uncritical, ~~particular~~ particularly in view of the fact that the interrupt handling configuration can be preloaded. The selection of the interrupt handling configuration to be preloaded can be carried out by the CT or in any other instance. It is also possible, to restrict the runtime of the configuration on the resource to which the 'interrupt processing has been assigned. ~~Reference is being had to~~ Regarding this, see PCT/DE 03/000942.

If the system has to react faster if an interrupt occurs, ~~it can be preferred to reserve in one embodiment,~~ a single resource, for example, a separate XPP-unit or parts of a data processing logic cell field, may be reserved for the execution of interrupt handling routines. In this case, it is also possible to preload interrupt handling routines for interrupts that are particularly critical. It is also possible to immediately start loading of an interrupt handling routine once the interrupt occurs. The selection of the configuration necessary for a respective interrupt, can be effected by triggering, wave-processing ~~and so forth., etc.~~

~~It is to be mentioned that by~~ By the methods described, it becomes possible to provide for an instantaneous reaction to the interrupt by using load-/store configurations in order to obtain a code-reentrancy. ~~Here, following~~ Following every single or every other data processing configuration, for example every five- or ten data processing configurations, a store configuration ~~is~~ may be executed and then a load configuration accessing the very memory arrays in which data have ~~been just been~~ written ~~is~~ may be carried out. Then, ~~it only has to be made sure that the memory areas used by that the store configuration remain untouched has to be ensured~~ until the configuration or group of configurations ~~that for which~~ the preloading has been effected ~~for~~ has been finished by completely executing a further store configuration. In ~~such a~~ this way of intermediately carried out load-/store configurations and simultaneous protection of not yet overaged store-memory areas, code-reentrancy is generated very easily, for example in compiling a program. Here, resource reservation may be advantageous as well.

Further, ~~a particular preferred in one example~~ in one example embodiment of the present invention, a reaction to an interrupt ~~consists in~~ may include using interrupt routines where code for the data processing logic cell field is forbidden, ~~this.~~ This embodiment ~~being preferred whets~~ may be particularly suited for an instance where one of the resources available is a sequential CPU.

In other words, an interrupt handling routine is executed only on a sequential CPU without calling data processing steps or routines making use of a data processing logic cell field. This ~~guarantees~~may guarantee that the processing on the data processing logic cell field is not ~~to be interrupted and then.~~ Then, further processing on ~~that~~the data processing logic cell field can be effected following a task switch. Although the actual interrupt routine does not ~~comprise~~include any data processing logic cell field code such as XPP-code, it can still be ~~made sure~~ensured that, at a later time no more relevant to real time processing capabilities, the data processing logic cell field reacts to an interrupt and/or a real time request determined, to state, information and/or data using the data processing logic cell field.

10

Third major part

~~Abstract—~~ Nowadays, the datapaths of modern microprocessors reach their limits by using static instruction sets. A way out of this limitations is a dynamic reconfigurable processor datapath extension achieved by integrating traditional static datapaths with the coarse grain dynamic reconfigurable XPP architecture (eXtreme Processing Platform). Therefore, a loosely asynchronous coupling mechanism of the corresponding datapath units has been developed and integrated onto a CMOS 0.13 μm standard cell technology from UMC. Here the SPARC compatible LEON processor is used, whereas its static pipelined instruction datapath has been extended to be configured and personalized for specific applications. This allows a various and efficient use, e.g. in streaming application domains like MPEG 4, digital filters, mobile communication modulation, etc. The chosen coupling technique allows asynchronous concurrency of the additionally configured compound instructions, which are integrated into the programming and compilation environment of the LEON processor.

6.1—Introduction

Fourth major part

1—Introduction

COMPILING AN HLL SUBSET EXTENDED BY PORT ACCESS FUNCTIONS TO AN RDFP

~~This document~~ The following describes a method, according to an embodiment of the present invention, for compiling a subset of a high-level programming language (HLL)-like, e.g., C or FORTRAN, extended by port access functions, to a reconfigurable data-flow processor (RDFP)-as described in Section 3. The program is may be transformed to a configuration of the RDFP.

This method can be used as part of an extended compiler for a hybrid architecture ~~consisting of~~ including a standard host processor and a reconfigurable data-flow coprocessor. The extended compiler handles a full HLL-like, e.g., standard ANSI C. It maps suitable program parts like, such as inner loops, to the coprocessor and the rest of the program to the host processor. It is also possible to map separate program parts to separate configurations.

However, these extensions are not the subject of this document. the discussion below.

2—Compilation Flow

~~This section briefly describes the phases of the~~ The compilation method. ~~may include a~~
frontend phase, a control/dataflow graph generation phase, and a configuration code phase.

2.1—Frontend

- 5 The compiler ~~uses~~ may use a standard frontend which translates the input program, (e.g., a C program) into an internal format ~~consisting of~~ including an abstract syntax tree (AST) and symbol tables. The frontend may also ~~performs~~ perform well-known compiler optimizations as, e.g., constant propagation, dead code elimination, common subexpression elimination ~~etc.~~ For details, ~~refer to any compiler construction textbook like [1]. The SURF compiler [2], etc.~~
- 10 For details regarding this, see A.V. Aho, R. Sethi, and J.D. Ullman, “Compilers – Principles, Techniques, and Tools,” Addison-Wesley 1986. The SURF compiler is an example of a compiler providing such a frontend. Regarding the SURF compiler, see The Stanford SUIF Compiler Group Homepage at <http://suif.stanford.edu>.

2.2—Control/Dataflow Graph Generation

15 Next, the program ~~is~~ may be mapped to a control/dataflow graph (CDFG) ~~consisting of~~ including connected RDFP functions. This phase is ~~the main subject of this document and presented in Section 4.~~ discussed in more detail below.

2.3—Configuration Code Generation

20 Finally, the last phase may directly ~~translate~~ translate the CDFG to configuration code used to program the RDFP. For PACT XPPT™ Cores, the configuration code ~~is~~ may be generated as an NML (Native Mapping Language) ~~file~~ file.

3—Configurable Objects and Functionality of an RDFP

- 25 ~~This section describes the configurable objects and functionality of a RDFP.~~ A possible implementation of the RDFP architecture is a PACT XPPT™ Core. Here we Discussed herein are only describe the minimum requirements for ~~a~~ an RDFP for this compilation method to work. The only data types considered are multi-bit words called data and single-bit control signals called events. Data and events are always processed as packets, ~~cf. Section 3.2.~~
- 30 See that which is discussed below under the heading “Packet-based Communication Network.” Event packets are called 1-events or 0-events, depending on their bit-value.

3.1—Configurable Objects and Functions

An RDFP ~~consists of~~includes an array of configurable objects and a communication network. Each object can be configured to perform certain functions ~~(, such as those listed below)~~. It ~~performs~~may perform the same function repeatedly until the configuration is changed. The array ~~needs~~need not be completely uniform, *i. e.*, not all objects need to be able to perform all functions. ~~E. g.~~For example, a RAM function can be implemented by a specialized RAM object ~~which~~that cannot perform any other functions. It is also possible to combine several objects to a “macro” to realize certain functions. ~~Several~~For example, several RAM objects can, ~~e. g.~~, be combined to ~~realize~~obtain a RAM function with larger storage.

~~The following~~Fig. 54 is a graphical representation of functions for processing data and event packets that can be configured into an RDFP. See Fig. 1 for a graphical representation. ~~The functions are as follows.~~

- ALU[opcode]: ALUs perform common ~~arithmetical~~arithmetic and —logical operations on data. ALU functions (“opcodes”) must be available for all operations used in the HLL.^[1] Otherwise, programs including operations that do not have ALU opcodes in the RDFP must be excluded from the supported HLL subset or substituted by “macros” of existing functions. ALU functions have two data inputs, A and B, and one data output, X. Comparators have an event output U instead of the data output. They produce a 1-event if the comparison is true, and a 0-event otherwise.
- CNT: ACNT is a counter function which has data inputs, LB, UB, and INC (lower bound, upper bound, and increment), and data output X (counter value). A packet at event input START starts the counter, and event input NEXT causes the generation of the next output value (and output events) or causes the counter to terminate if UB is reached. If NEXT is not connected, the counter ~~counts~~may count continuously. The output events U, V, and W have the following functionality: For a counter counting N times, N-1 0-events and one 1-event aremay be generated at output U. At output V, N 0-events aremay be generated, and at output W, N 0-events and one 1-event aremay be created. The 1-event

¹ ~~Otherwise programs containing operations which do not have ALU opcodes in the RDFP must be excluded from the supported HLL subset or substituted by “macros” of existing functions.~~

at W is only created after the counter has terminated, *i.e.*, a NEXT event packet was received after the last data packet was output.

- 5 • RAM[size]: The RAM function ~~stores~~may store a fixed number of data words (“size”). It has a data input RD and a data output OUT for reading at address RD. Event output ERD signals completion of the read access. For a write access, data inputs WR and IN (address and value) and data output OUT ~~is~~may be used. Event output EWR signals completion of the write access. ERD and EWR always generate 0-events. Note that external RAM can be handled as RAM functions exactly like internal RAM.
- 10 • GATE: A GATE ~~synchronizes~~may synchronize a data packet at input A ~~back~~ and an event packet at input E. When both inputs have arrived, they ~~are~~may both be consumed. The data packet ~~is~~may be copied to ~~the~~ output X, and the event packet to output U.
- 15 • MUX: An MUX function ~~has~~may have 2 data inputs, A and B, an event input, SEL, and a data output, X. If SEL receives a 0-event, input A ~~is~~may be copied to output X ~~and input B, and input B may be~~ discarded. For a 1-event, B ~~is~~may be copied, and A may be discarded.
- 20 • MERGE: A MERGE function ~~has~~may have 2 data inputs, A and B, an event input SEL, and a data output X. If SEL receives a 0-event, input A ~~is~~may be copied to output X, but input B is *not* discarded. The packet ~~is~~may be left at the input B instead. For a 1-event, B ~~is~~may be copied and A left at the input.
- 25 • DEMUX: A DEMUX function ~~has~~may have one data input A, an event input SEL, and two data outputs X and Y. If SEL receives a 0-event, input A ~~is~~may be copied to output X, and no packet is created at ~~the~~ output Y. For a 1-event, A may be copied to Y, and no packet is copied to Y, and no packet is created at output X.
- 30 • MDATA: A MDATA function ~~multiplies~~may multiply data packets. It ~~has~~may have a data input A, an event input SEL, and a data output X. If SEL receives a 1-event, a data packet at A ~~is~~may be consumed and copied to output X. For all subsequent 0-~~event~~events at SEL, a copy of the input data packet ~~is~~may be produced at the output

without consuming new packets at A. Only if another 1-event arrives at SEL, the next data packet at A is may be consumed and copied.^[2]It is noted that this can be implemented by a MERGE with special properties on XPP™.

- 5 • INPORT[name]: ~~Receives~~An INPORT function may receive data packets from— outside the RDFP through input port “name” and ~~copies~~may copy them to data output X. If a packet was received, a 0-event is may be produced at event output U, too. (~~Note~~It is noted that this function can only be configured at special objects connected to external busses.)

- 10 • OUTPORT[name]: ~~Sends~~An OUTPORT function may send data packets received at data input A to the outside of the RDFP through output port “name.”. If a packet was sent, a 0-event is may be produced at event output U, too. (~~Note~~It is noted that this function can only be configured at special objects connected to external busses.)

15 Additionally, the following functions manipulate only event packets:

- 0-FILTER, 1-FILTER: A FILTER ~~has~~may have an input E and an output U. A 0-FILTER ~~copies~~may copy a 0-event from E to U, but 1-EVENTs at E are discarded. A 1-FILTER ~~copies~~may copy 1-events and ~~discards~~discard 0-events.

20

- INVERTER: ~~Copies~~An INVERTER may copy all events from input E to output U, but ~~inverts its value.~~invert their values.

- 0-CONSTANT, 1-CONSTANT: 0-CONSTANT ~~copies all events from input E to output U but changes~~may copy all events from input E to output U, but may change them all to value 0. 1-CONSTANT ~~changes~~may change them all to value 1.

25

- ECOMB: ~~Combines~~ECOMB may combine two or more inputs E1, E2, E3 . . . , producing a packet at output U. The output is may be a 1-event if and only if one or more of the input packets are 1-events (logical *or*). A packet must be available at all inputs before an

30

² ——— ~~Note that this can be implemented by a MERGE with special properties on XPP™.~~

output packet is produced.^{[[3]]} It is noted that this function may be implemented by the EAND operator on the XPP™.

- ESEQ[seq]: An ESEQ ~~generates~~may generate a sequence “seq” of events, e.g., “0001”,0001, at its output U. If it has an input START, one entire sequence ~~is~~may be generated for each event packet arriving at U. The sequence is only repeated if the next event arrives at U. However, if START is not connected, ESEQ may constantly ~~repeats~~repeat the sequence.

10 ~~Note~~It is noted that the ALU, MUX, DEMUX, GATE and ECOMB functions may behave like their equivalents in ~~elassical dataflow machines [3, 4].~~ conventional dataflow machines. In this regard, see A.H. Veen, “Dataflow Architecture,” *ACM Computing Surveys*, 18(4) (December 1986); and S.J. Allan & A.E. Oldehoeft, “A Flow Analysis Procedure for the Translation of High-Level Languages to a Data Flow Language,” *IEEE Transactions on*
15 *Computers*, C-29(9):826-831 (September 1980).

3.2—Packet-based Communication Network

The communication network of an RDFP can connect ~~an~~ outputs of one object, (*i. e.*, its respective function), to the input(s) of one or several other objects. This is usually achieved
20 by busses and switches. By placing the functions properly on the objects, many functions can be connected arbitrarily up to a limit imposed by the device size. As mentioned above, all values ~~are~~may be communicated as packets. A separate communication network ~~exists~~may exist for data and event packets. The packets may synchronize the functions as in a dataflow machine with acknowledge [3]. ~~I.e.,~~ In this regard, see A.H. Veen, *supra*. That is, the
25 function only executes when all input packets are ~~avail-able~~available (apart from the non-strict exceptions as described above). The function may also ~~stalls~~stall if the last output packet has not been consumed. Therefore, a data-flow graph mapped to an RDFP may self-synchronize ~~synchronizes~~synchronize its execution without the need for external control. Only if two or
30 more function, outputs (data or event) are connected to the same function input (“N to 1-connection”), is the self-synchronization ~~is disabled~~.^{[[4]]} It is noted that on XPP™ Cores, an “N to 1 connection” for events is realized by the EOR function, and, for data, by just

³ Note that this function is implemented by the EAND operator on the XPP™.

assigning several outputs to an input. The user has to ensure that only one packet arrives at a time in a correct CDFG. Otherwise, a packet might get lost, and the value resulting from combining two or more packets is- undefined. However, a function output can be connected to many function inputs (“1 to N connection”) without problems.

5

There are some special cases:

- A function input can be *preloaded* with a distinct value during configuration. This packet ~~is~~may be consumed like a normal packet coming from another object.
- A function input can be defined as *constant*. In this case, the packet at the input ~~is~~may be reproduced repeatedly for each function execution.

10

An RDFP ~~requires~~may require register delays in the dataflow. Otherwise, very long combinational delays and asynchronous feedback is possible. ~~We assume~~It is assumed that delays are inserted at the inputs of some functions (like for most ALUs) and in some routing segments of the communication network. ~~Note~~It is noted that registers may change the tuning, but not the functionality, of a correct CDFG.

15

4——Configuration Generation

4.1——Language Definition

20

The following HLL features are not supported by the method described ~~here~~herein:

- pointer operations
- library calls, operating system calls (including standard I/O functions)
- recursive function calls (~~Note that~~ non-recursive function calls can be eliminated by function in-lining and therefore are not considered ~~here~~herein.)
- All scalar data types ~~are~~may be converted to type integer. Integer values ~~are~~may be equivalent to *data* packets in the— RDFP. Arrays (~~Possibly~~possibly multi-dimensional) are the only composite data types considered.

25

The following additional features are supported:

⁴——~~Note that on XPP™ Cores, a “N to I connection” for events is realized by the FOR function, and for data by just assigning several outputs to an input.~~

IMPORTS and OUTPUTS can be accessed by the HLL functions *getstream(name, value)* and *putstream(name, value)*, respectively.

4.2 ~~Mapping of High-Level Language Constructs~~

5 This method ~~converts~~ may convert an HLL program to a CDFG ~~consisting of~~ including the RDFP functions defined in ~~Section 3.1~~ the discussion under the heading “Configurable Objects and Functions.” Before the processing starts, all HLL program arrays ~~are~~ may be mapped to RDFP RAM functions. An array *x* ~~is~~ may be mapped to RAM *RAM(x)*. If several arrays are mapped to the same RAM, an offset ~~is~~ may be assigned, too. The RAMs ~~are~~ may be added to an initially empty CDFG. There must be enough RAMs of sufficient size for all
10 program arrays.

The CDFG ~~is~~ may be generated by a traversal of the AST of the HLL program. It ~~processes~~ may process the program statement by statement and ~~descends~~ descend into the.
15 loops and conditional statements as appropriate. The following two pieces of information ~~are~~ may be updated at every program point^[5], (which refers to a point between two statements or before the beginning or after the end of a program component such as a loop or a conditional statement), during the traversal:

- START ~~points~~ may point to an event output of ~~an~~ RDFP function. This output
20 ~~delivers~~ may deliver a 0-event whenever the program execution reaches this program point. At the beginning, a 0-CONSTANT preloaded with an event input ~~is~~ may be added to the CDFG. (It ~~delivers~~ may deliver a 0-event immediately after configuration.))
START may initially ~~points~~ point to its output. This event ~~is~~ may be used to start the overall program execution. ~~The~~ A *START_{new}* signal generated after a program part has
25 finished executing ~~is~~ may be used as new START signal for the following program parts, or it ~~signals~~ may signal termination of the entire program. The START events may guarantee that the execution order of the original program is maintained wherever the data dependencies alone are not sufficient. This scheduling scheme ~~is~~ may be similar to a *one-hot controller* for digital hardware.
- VARLIST ~~is~~ may be a list of {*variable, function-output*} pairs. The pairs may map
30 integer variables or array elements to a CDFG function’s output. The first pair for a

⁵ ~~———— In a program, program points are between two statements or before the beginning or after the end of a program component like a loop or a conditional statement.~~

variable in VARLIST ~~contains~~may contain the output of the function which produces the value of this variable valid at the current program point. New pairs ~~are~~may be always added to the front of VARLIST. The expression VARDEF(var) ~~refers~~may refer to the *function-output* of the first pair with *variable* var in VARLIST.^[6] ~~With respect to this way of using a VARLIST, see D. Galloway, "The transmogrifier C hardware description language and compiler for FPGAs," Proc. FPGAs for Custom Computing Machines, IEEE Computer Society Press, 1995, at 136-44.~~

~~The following subsections~~Below are systematically list all HLL listed HLL program components and ~~describe~~descriptions of how they ~~are~~may be processed, thereby altering the CDFG, START, and VARLIST.

4.2.1—Integer Expressions and Assignments

Straight-line code without array accesses can be directly mapped to a data-flow graph. One ALU ~~is~~may be allocated for each operator in the program. Because of the self-synchronization of the ALUs, no explicit control or scheduling is needed. Therefore processing these assignments does not access or alter START. The data ~~dependences~~dependencies (as they would be exposed in the DAG representation of the program^[1]) ~~are, in regard to which see A.V. Aho et al., supra) may be~~ analyzed through the processing of VARLIST. These assignments may synchronize themselves through the data-flow. The data-driven execution may automatically ~~exploit~~exploit the available instruction level parallelism.

All assignments may evaluate the right-hand side (RHS) or source expression. This evaluation ~~results~~may result in a pointer to a CDFG object's output (or pseudo-object as defined below). For integer assignments, the left-hand side (LHS) variable or destination ~~is~~may be combined with the RHS result object to form a new pair {LHS, result(RHS)} which ~~is~~may be added to the front of VARLIST.

For the following examples, C syntax is used. The simplest statement ~~is~~may be a constant assigned to an integer:^[7]

⁶ ——— This method of using a VARLIST is adapted from the Transmogrifier C compiler [5].

⁷ ——— Note that we use C syntax for the following examples.

a = 5;

It ~~doesn't~~does not change the CDFG, but adds {a, 5} to the front of VARLIST. —The constant 5 is —a “pseudo-object” which only holds the value,— but —does not refer to a CDFG object. Now VARDEF(a) equals 5 at subsequent program points before a is

5 redefined.

Integer assignments can also combine variables already defined and constants:

b = a * 2 + 3;

In the AST, the RHS is already converted to an expression tree. This tree ~~is~~may be transformed to a combination of old and new CDFG objects (which are added to the CDFG) as follows:—Each operator (internal node) of the tree ismay be substituted by an ALU with the opcode corresponding to the operator in the tree. If a leaf node is a constant, the ALU's input ismay be directly connected to that constant. —If a leaf node is an integer variable var, it ismay be looked up in VARLIST, *i.e.* —VARDEF(var) is retrieved. Then VARDEF(var) (an

15 output of an already existing object in CDFG or a constant) ismay be connected to the ALU's input. The output of the ALU corresponding- to the root operator in the expression tree is defined as the *result* of the RHS. Finally, a new pair {LHS, result(RHS)} ismay be added to VARLIST. If the two assignments above are processed, the CDFG with two ALUs, as shown in Fig.—2 ~~is created.~~⁸ 55, may be created. It is noted that the input and output names can be

20 deduced from their position. It is futher noted that the compiler frontend would normally have substituted the second assignment by b = 13 (constant propagation). For the simplicity, no frontend optimizations are considered in this and the following examples. Outputs occurring in VARLIST are labeled by Roman numbers. After these two assignments, VARLIST = [{b, I}, {a, 5}]. (The front of the list is on the left side.) Note that all inputs

25 connected to a constant (whether direct from the expression tree or retrieved from VARLIST) must be defined as constant. —Inputs defined as constants have a small c next to the input arrow in Fig.—2.—55.

4.2.2—Conditional Integer Assignments

⁸ —Note that the input and output names can be deduced from their position, cf. Fig. 1. Also note that the compiler front-end would normally have substituted, the second assignment by b = 3.3 (constant propagation). For the simplicity of this explanation, no frontend optimizations are considered in this and the following examples.

For conditional if-then-else statements ~~containing~~including only integer assignments, objects for condition evaluation ~~are~~may be created first. The object event output indicating the condition result ~~is~~may be kept for choosing the correct branch result later. Next, both branches ~~are~~may be processed in parallel, using separate copies VARLIST1 and VARLIST2 of VARLIST. (VARLIST itself is not changed.) Finally, for all variables added to VARLIST1 *or* VARLIST2, a new entry for VARLIST ~~is~~may be created (combination phase). The valid definitions from VARLIST1 and VARLIST2 ~~are~~may be combined with a MUX function, and the correct input ~~is~~may be selected by the condition result. For variables only defined in one of the two branches, the multiplexer ~~uses~~may use the result retrieved from the original VARLIST for the other branch. If the original VARLIST does not have an entry for this variable, a special “undefined” constant value ~~is~~may be used. However, in a ~~function-~~functionally correct program, this value will never be used. As an optimization, only variables *live* ~~{1}~~^[9](see A.V. Aho et al., *supra*) after the if-then-else structure need to be added to VARLIST in the combination phase.^[9]A variable is *live* at a program point if its value is read at a statement reachable from the point without intermediate redefinition.

Consider the above with respect to the following example:

```

i = 7;
a = 3;
20  if (i < 10) {
        a = 5;
        c = 7;
    }
    else {
25      c = a - 1;
        d = 0;
    }

```

For this example, Fig. ~~3~~56 shows the resulting CDFG. Before the if-then-else construct, VARLIST = [{a, 3}, {i, 7}]. After processing the branches, for the then branch, VARLIST1

⁹ ~~Definition: A variable is live at a program point if its value is read at a statement reachable from here without intermediate redefinition.~~

= [{ c, 7}, {a, 5}, {a, 3}, {i, 7}], and for the else branch, VARLIS72 = [{d, 0}, {c, I }, {a, 3}, {i, 7 }]. After combination, VARLIST = [{d, II }, {c, III}, {a, IV}, {a, 3}, {i, 7}].

Note that case- or switch-statements can be processed, too, since they can ~~—be converted,~~
without loss of generality ~~—be converted,~~ to nested if-then-else statements.

Processing conditional statements this way does not require explicit control and does not change START. Both branches ~~are~~may be executed in parallel and synchronized by the ~~data-~~dataflow. It is possible to pipeline the dataflow for optimal throughput.

4.2.3—General Conditional Statements

Conditional statements ~~containing~~including either array accesses (ef. ~~Section 4. 2. 7~~
~~below~~see the discussion below under the heading “Array Accesses”) or inner loops cannot
~~the~~be processed as described in ~~Section 4. 2. 2.~~above under the heading “Conditional Integer
Assignments.” Data packets must ~~only~~only be sent only to the active branch. This ~~is~~may be
achieved by the implementation shown in Fig. ~~8,~~Fig. 57, similar to the method presented in
[4]. ~~S.J. Allan et al., supra.~~

A dataflow analysis ~~is~~may be performed to compute *used sets* use and *defined sets* def ~~[1]~~of
~~both branches.~~¹⁰ (see A.V. Aho et al., supra) of both branches. A variable is used in a
statement (and hence in a program region including the statement) if its value is read. A
variable is defined in a statement (or region) if a new value is assigned to it. For the current
VARLIST entries of all variables in $IN = use(thenbody) \cup def(thenbody) \cup use(elsebody)$
 $\cup def(elsebody) \cup use(header)$, DEMUX functions controlled by the IF condition are
inserted. ~~Note~~It is noted that arrows with double lines in Fig. ~~8~~57 denote connections for all
variables in IN, and the ~~shaded~~shadowed DEMUX function stands for several DEMUX
functions, one for each variable in IN. The DE- MUX functions forward data packets only to
the selected branch. New lists VARLIST1 and VARLIST2 are compiled with the respective
outputs of these DEMUX functions. The then-branch is processed with VARLIST1, and the
else branch with VARLIST2. Finally, the output values are combined. OUT
~~contains~~includes the new values for the same variables as in IN. Since only one branch is

¹⁰ ~~A variable is used in a statement (and hence in a program region containing this statement) if its value is read. A variable is defined in a statement (or region) if a new value is assigned to it.~~

ever activated, there will not be a conflict due to two packets arriving simultaneously. The combinations will be added to VARLETVARLIST after the conditional statement. If the IF execution shall be pipelined, MERGE opcodes for the output must be inserted, too. They are controlled by the condition like the DEMUX functions.

5

~~The following extension with~~ With respect to [4] ~~is added (dotted lines in Fir. 8) that which is dicussed in S.J. Allan et al., *supra*, the following extension, corresponding to the dashed lines of Fig. 57 may be added in an embodiment of the present invention in order to control the execution as mentioned above with START events:-.~~ The START input is ~~may be~~ ECOMB-
 10 combined with the condition output and connected to the SEL input of the DEMUX
functions. The START inputs of thenbody and elsebody ~~are~~ may be ~~generated from the~~
ECOMB output sent through a 1-FILTER and a 0-CONSTANT[[¹¹]] or through a 0-FILTER,
respectively. (The 0-CONSTANT may be required since START events must always be 0-
events.) The overall $START_{new}$ output is ~~may be~~ generated by a simple “2 to 1 connection” of
 15 thenbody’s and elsebody’s $START_{new}$ outputs. With this extension, arbitrarily nested
conditional statements or loops can be handled within thenbody and elsebody.

4.2.4—WHILE Loops

WHILE loops ~~are~~ may be processed similarly to the scheme presented in [4], cf. Fig. 9. ~~As~~
 20 ~~in Section 4. 23, double~~ S.J. Allan et al., *supra* (see Fig. 58). Double line connections and
~~shaded~~ shadowed MERGE and DEMUX functions represent duplication for all variables in
 IN. Here $IN = use(whilebody) \cup def(whilebody) \cup use(header)$. The WHILE loop
~~executes~~ may execute as follows:-. In the first loop iteration, the MERGE functions ~~may~~
select all input values from VARLIST at loop entry (SEL=0). The MERGE outputs ~~are~~ may
 25 be connected to the header and the DEMUX functions. If the while condition is true (SEL=1),
the input values ~~are~~ may be forwarded to the whilebody, and otherwise to OUT. The output
values of the while body ~~are~~ may be fed back to whilebody’s input via the MERGE and
DEMUX operators as long as the condition is true. Finally, after the last iteration, they
are ~~may be~~ forwarded to OUT. The outputs ~~are~~ may be added to the new VARLIST. [[¹²]] It is

¹¹ ——— The 0-CONSTANT is required since START events must always be 0.events.

¹² ——— Note that the MERGE function for variables not live at the loop’s beginning and the whilebody’s beginning can be removed since its output is not used. For these variables, only the DEMUX function to output the final value is required. Also note that the MERGE functions can be replaced by simple “2-to-1 connections” if the configuration process guarantees that packets from IN1 always arrive at the DEMUX’s input before feedback values arrive.

noted that the MERGE function for variables not live at the loop's beginning and the whilebody's beginning can be removed since its output is not used. For these variables, only the DEMUX function to output the final value is required. It is further noted that the MERGE functions can be replaced by simple "2 to 1 connections" if the configuration process guarantees that packets from IN1 always arrive at the DEMUX's input before feedback values arrive.

~~Two extensions with respect to [4] are added (dotted lines in Fir. 9):~~

With respect to that which is dicussed in S.J. Allan et al., *supra*, the following two extensions, corresponding to the dashed lines in Fig. 58, may be added in an embodiment of the present invention.

- In [4]S.J. Allan et al., *supra*, the SEL— input of the MERGE functions is preloaded with 0. ~~Hence~~Thus, the loop execution begins immediately and can be executed only once. Instead, ~~we connect~~in an embodiment of the present invention, the START input may be connected to the MERGE's SEL input ("2 to 1 connection" with the header output). This ~~allows to~~may allow control of the time of the start of the loop execution and ~~to~~may allow its restart-it.
- The whilebody's START input ~~is~~may be connected to the header output, sent through a 1-FILTER/0-CONSTANT combination as above (generates a 0-event for each loop iteration). By ECOMB-combining- whilebody's $START_{new}$ output with the header output for the MERGE functions' SEL inputs, the next loop iteration is only started after the previous one has finished. The while loop's $START_{new}$ output is generated by filtering the header output for a 0-event.

With these extensions, arbitrarily nested conditional statements or loops can be handled within whilebody.

4.2.5—FOR Loops

FOR loops are particularly regular WHILE loops. Therefore ~~we could handle them~~, they may handled as explained above. However, ~~our RDFP features the~~an RDFP according to an embodiment of the present invantion may feature a special counter function CNT and ~~the~~a

data packet multiplication function MDATA₁ which can be used for a more efficient implementation of FOR loops. This new FOR loop scheme is shown in Fig. 10-59.

A FOR loop ~~is~~may be controlled by a counter CNT. The lower bound (LB), upper bound (UB), and increment (INC) expressions ~~are evaluated like any other expressions (see Sections 4. 2. 1 and 4. 2. 7)~~may be evaluated like any other expression (see, for example, that which is discussed above under the heading “Integer Expressions and Assignments,” and that which is discussed below under the heading “Array Accesses”) and connected to the respective inputs.

As opposed to WHILE loops, a MERGE/DEMUX combination is only required for variables in $IN1 = \text{def}(\text{forbody})$, i. e., those *defined* in forbody. ~~[[¹³]]-IN1 does not contain~~It is noted that the MERGE functions can be replaced by simple “2 to 1 connections” as for WHILE loops if the configuration process guarantees that packets from IN1 always arrive at the DEMUX’s input before feedback values arrive. IN1 does not include variables which are only *used* in forbody, LB, UB, or INC, and ~~does also~~does not contain the loop index variable. Variables in IN1 ~~are~~may be processed as in WHILE loops, but the MERGE and DEMUX functions’ SEL input is connected to CNT’s W output. (The W output ~~does~~may do the inverse of a WHILE loop’s header output; ~~it outputs a 1.~~It may output a 1-event after the counter has terminated. Therefore, the inputs of the MERGE functions and the outputs of the DEMUX functions ~~are~~may be swapped here, and the MERGE functions’ SEL inputs ~~are~~may be preloaded with 1-events.)

CNT’s X output ~~provides~~may provide the current value of the loop index variable. If the final index value is required (live) after the FOR loop, it ~~is~~may be selected with a DEMUX function controlled by CNT’s U event output (which ~~produces~~may produce one event for every loop iteration).

Variables in $IN2 = \text{use}(\text{forbody}) \setminus \text{def}(\text{forbody})$, i. e., those defined outside the loop and only used (but not redefined) inside the loop ~~are~~may be handled differently. Unless it is a constant value, the variable’s input value (from VARLIST) must be reproduced ~~in~~ each loop

iteration since it is consumed in each iteration. Otherwise, the loop would stall from the second iteration onwards. The packets ~~are~~may be reproduced by MDATA functions, with the SEL inputs connected to CNT's U output. The SEL inputs must be preloaded with a 1-event to select the first input. The 1-event provided by the last iteration ~~selects~~may select a new value for the next execution of the entire loop.

The following control events (~~corresponding to the~~ dotted lines in Fig. ~~40 59~~) are similar to the WHILE loop extensions, but simpler. CNT's START input ~~is~~may be connected to the loop's overall START signal. $START_{new}$ ~~is~~may be generated from CNT's W output, sent through a 1-FILTER and 0-CONSTANT. CNT's V output ~~produces~~may produce one 0-event for each loop iteration and ~~is~~may therefore ~~be~~ used as forbody's START. Finally, CNT's NEXT input ~~is~~may be connected to forbody's $START_{new}$ output.

For pipelined loops (as defined below ~~in Section 4. 2. 6~~); under the heading "Vectorization and Pipelining"), loop iterations ~~are~~may be allowed to overlap. Therefore, CNT's NEXT input ~~needs~~need not be connected. Now the counter ~~produces~~may produce index variable values and control events as fast as they can be consumed. However, in this case CNT's W output is not sufficient as overall $START_{new}$ output since the counter terminates before the last iteration's forbody finishes. Instead, $START_{new}$ ~~is~~may be generated from CNT's U output ECOMB-combined with forbody's $START_{new}$ output, sent through a 1-FILTER/0-CONSTANT combination. The ECOMB ~~produces~~may produce an event after termination of each loop iteration, but only the *last* event is a 1-event because only the last output of CNT's U output is a 1-event. ~~Hence~~Thus, this event ~~indicates~~may indicate that the last iteration has finished. ~~Cf. Section 4. 3 for a~~ FOR loop example compilation with and without pipelining— is provided below under the heading "More Examples."

As for WHILE loops, these methods allow ~~to process~~for arbitrarily processing nested loops and conditional statements. The following advantages over WHILE loop implementations ~~are~~may be achieved:

¹³ ——— Note that the MERGE functions can be replaced by simple "2-to-1 connections" as for WHILE loops if the configuration process guarantees that packets from IN1 always arrive at the DEMUX's input before feedback values arrive.

- One index variable value ~~is~~may be generated by the CNT function each clock cycle. This is faster and smaller than the WHILE loop implementation which allocates a MERGE/DEMUX/ADD loop and a comparator for the counter functionality.

- 5 • Variables in IN2 (only used in forbody) ~~are~~may be reproduced in the special MDATA functions and need not go through a MERGE/DEMUX loop. This is again faster and smaller than the WHILE loop implementation.

4.2.6—Vectorization and Pipelining

10 ~~The method~~In the embodiments described so far generates above, CDFGs ~~Performing are~~generated that perform the HLL program's functionality on an RDFP. However, the program execution is unduly sequentialized by the START signals. In some cases, innermost- loops can be *vectorized*. This means that loop iterations can overlap, leading to a pipelined dataflow through the operators of the loop body. The *Pipeline Vectorization* technique {6}

15 ~~can~~(see Markus Weinhardt et al., "Pipeline Vectorization," *supra*) can be easily applied to the compilation method ~~presented hereof~~embodiments of the present invention. As mentioned above, for FOR loops, the CNT's NEXT input ~~is~~may be removed so that CNT counts continuously, thereby overlapping the loop iterations.

20 All loops without array accesses can be pipelined since the dataflow automatically synchronizes ~~loop-carried dependences~~, i.e. ~~dependences~~dependencies, i.e., dependencies between a statement in one iteration and another statement in a subsequent iteration. Loops with array accesses can be ~~pipelined~~ if the array, (i. e., RAM), accesses do not cause loop-carried ~~dependences~~dependencies or can be transformed to such a form. In this case, no

25 RAM address is written in one iteration and read in a subsequent iteration. Therefore, the read and write accesses to the same RAM may overlap. This degree of freedom is exploited in the RAM access technique described below. Especially for dual-ported RAM, it leads to considerable performance improvements.

30 4.2.7—Array Accesses

In contrast to scalar variables, array accesses have to be controlled explicitly in order to maintain the program's correct execution order. As opposed to normal dataflow machine models {3}, ~~a~~(see A.H. Veen, *supra*), an RDFP does not have a single address space. Instead,

the arrays ~~are~~may be allocated to several RAMs. This leads to a different approach to handling RAM accesses and opens up new opportunities for optimization.

To reduce the complexity of the compilation process, array accesses ~~are~~may be processed in two phases. Phase 1 ~~uses~~may use “pseudo-functions” for RAM read and write accesses. A RAM read function ~~has~~may have an RD data input (read address) and an OUT data output (read value), and a RAM write function ~~has~~may have WR and IN data inputs (write address and write value). Both functions are labeled with the array the access refers to, and both may have a START event input and a U event output. The events may control the access order.

In Phase 2, all accesses to the same RAM ~~are~~may be combined and substituted by a single RAM function ~~as shown in Fig. 1.~~ This ~~involves~~may involve manipulating the data and event inputs and outputs such that the correct execution order is maintained and the outputs are forwarded to the correct part of the CDFG.

● Phase 1:

Phase 1—Since arrays ~~are~~may be allocated to several RAMs, only accesses to the same RAM have to be synchronized. Accesses to different RAMs can occur concurrently or even out of order. In case of data dependencies, the accesses may self-synchronize automatically. Within pipelined loops, not even read and write accesses to the same RAM have to be synchronized. This ~~is~~may be achieved by maintaining separate START signals for every RAM or even separate START signals for RAM read and RAM write accesses in pipelined loops. At the end of a basic block [1]¹⁴, all ~~START_{new}~~, which is a program part with a single entry and a single exit point, i.e., a piece of straight-line code, (see A.V. Aho et al., *supra*), all ~~START_{new}~~ outputs must be combined by a ~~an~~ ECOMB to provide a START signal for the next basic block, which guarantees that all array accesses in the previous basic block are completed. For pipelined loops, this condition can even be relaxed. Only after the loop exit, all accesses have to be completed. The individual loop iterations need not be synchronized.

First the RAM addresses ~~are~~may be computed. The compiler frontend’s standard transformation for array accesses can be used, and a CDFG function’s output ~~is~~may be

¹⁴ ——— A basic block is a program part with a single entry and a single exit point, i.e. a piece of straight-line code.

generated which ~~provides~~may provide the address. If applicable, the offset with respect to the RDFP RAM (as determined in the initial mapping phase) must be added. This output ~~is~~may be connected to the pseudo RAM read's RD input (for a read access) or to the pseudo RAM write's WR input (for a write access). Additionally, the OUT output (read) or IN input (write) ~~is~~may be connected. The START input ~~is~~may be connected to the variable's START signal, and the U output ~~is~~may be used as $START_{new}$ for the next access.

To avoid redundant read accesses, RAM reads ~~are~~may also be registered in VARLIST. Instead of an integer variable, an array element ~~is~~may be used as the first element of the pair. However, a change in a variable occurring in an array index invalidates the information in VARLIST. It must then be removed from it.

The following example with two read accesses compiles to the intermediate CDFG shown in Fig.-42-~~60~~. The START signals refer only to variable a. STOP1 is the event connection which synchronizes the accesses. Inputs START (old), i_1 and j should be substituted by the actual outputs resulting from the program before the array reads.

```
x = a[i];
y = a[j];
z = x+y;
```

Fig.-43-~~61~~ shows the translation of the ~~following~~ write access $[[[:]] a[i] = x;$

- Phase 2:**

~~Phase 2~~ ~~We now merge the~~ The pseudo-functions of all accesses may be merged to the same RAM and may be substitute ~~them~~ by a single RAM function. For all data inputs (RD for read access and WR and IN for write access), ~~GATEs are~~GATEs may be inserted between the input and the RAM function. Their E inputs ~~are~~may be connected to the respective START inputs of the original pseudo-functions. If a RAM is read and written at only one program point, the U output of the- read and write access ~~is~~may be moved to the ERD or EWR output, respectively. For example, the single access $a[i] = x;$ from Fig.-43-~~is 61~~ may be transformed to the final CDFG shown in Fig.-5-~~62~~.

However, if several read or several write accesses, (*i. e.*, pseudo-functions from different program points) to the same RAM occur, the ERD or EWR events are not specific anymore. But a *START_{new}* event of the original pseudo function should only be generated for the respective program point, *i.e.*, for the ~~current~~current access. This ~~is~~may be achieved by connecting the START signals of all *other* accesses (pseudo-functions) of the same type (read or write) with the *inverted* START signal of the current access. The resulting signal ~~produces~~may produce an event for every access, but ~~a 1-event for only for the current access~~a 1-event. This event ~~is~~may be ECOMB-combined with the RAM's ERD or EWR output. The ECOMB's output will only occur after the access is completed. Because ECOMB OR-combines its event packets, only the current access produces a 1-event. Next, this event ~~is~~may be filtered with a 1-FILTER and changed by a 0-CONSTANT, resulting in a *START_{new}* signal which produces a 0-event only after the current access is completed as required.

For several accesses, several sources ~~are~~may be connected to the RD, WR, and IN inputs of a RAM. This ~~disables~~may disable the self-synchronization. However, since only one access occurs at a time, the GATEs only allow one data packet to arrive at the inputs.

For read accesses, the packets at the OUT output face the same problem as the ERD event packets: ~~They, which is that they~~ occur for every read access, but must ~~only~~-be used (and forwarded to subsequent operators) only for the current access. This can be achieved by connecting the OUT output via a DEMUX function. The Y output of the DEMUX ~~is~~may be used, and the X output ~~is~~may be left unconnected. Then it ~~acts~~may act as a selective gate which only forwards packets if its SEL input receives a 1-event, and discards its data input if SEL receives a 0-event. The signal created by the ECOMB described above for the *START_{new}* signal ~~creates~~may create a 1-event for the current access, and a 0-event otherwise. Using it as the SEL input achieves exactly the desired functionality.

Fig. [[4]] 63 shows the resulting CDFG for the first example above (two read accesses), after applying the transformations of Phase 2 to Fig. ~~12.~~ 60. STOP1 ~~is now~~may be generated as follows: ~~START(old)~~ START(old) ~~is~~may be inverted, "2 to 1 connected" to STOP 1 (because it is the START input of the second read pseudo-function), ECOMB-combined with RAM's ERD output and sent through the 1-FILTER/0-CONSTANT combination. *START(new)* ~~is~~may be generated similarly, but here *START(old)* ~~is~~may be directly used and STOP1 inverted. The GATEs for input IN (i and j) ~~are~~may be connected to *START(old)* and STOP1, respectively,

and the DEMUX functions for outputs x and y ~~are~~may be connected to the ECOMB outputs related to STOP1 and START(new).

Multiple write accesses may use the same control events, but instead of one GATE per access for the RD inputs, one GATE for WR and one gate for IN (with the same E input) ~~are~~may be used. The EWR output ~~is~~may be processed like the ERD output for read accesses.

This transformation ~~ensures~~may ensure that all RAM accesses are executed correctly, but it is not very fast since read or write accesses to the same RAM are not pipelined. The next access only starts after the previous one is completed, even if the RAM being- used has several pipeline stages. This inefficiency can be removed as follows:-

First, continuous *sequences* of either read accesses or write accesses (not mixed) within a basic block ~~are~~may be detected by checking for pseudo-functions whose U output is directly connected to the START input of another pseudo-function of the same RAM and the same type (read or write). For these sequences, it is possible to stream data into the RAM rather than waiting for the previous access to complete. For this purpose, a combination of MERGE functions ~~selects~~may select the RD or WR and IN inputs in the order given by the sequence. The MERGES must be controlled by iterative ESEQs guaranteeing that the inputs are only forwarded in the desired order. Then only the first access in the sequence needs to be controlled by a GATE or GATEs. Similarly, the OUT outputs of a read access can be distributed more efficiently for a sequence. A combination of DEMUX functions with the same ESEQ control can be used. It ~~is~~may be most efficient to arrange the MERGE and DEMUX functions as balanced binary trees.

The $START_{new}$ signal ~~is~~may be generated as follows:- For a sequence of length n, the START signal of the entire sequence ~~is~~may be replicated n times by an ESEQ[00..1] function with the START input connected to the sequence's START. Its output ~~is~~may be directly "N to 1 connected" with the other accesses' START signal (for single accesses) or ESEQ outputs sent through 0-CONSTANT (for access sequences), ECOMB-connected to EWR or ERD, respectively, and sent through a 1-FILTER/0-CONSTANT combination, similar to the basic method described above. Since only the last ESEQ output is a 1-event, only the *last* RAM access generates a $START_{new}$ as required. Alternatively, for read accesses, the generation of

the last output can be sent through a GATE (without the E input connected), thereby producing a $START_{new}$ event.

Fig. 14 64 shows the optimized version of the first example (Figs. 60 and 63) using the ESEQ-method for generating $START_{new}$ and Fig. 6 65 shows the final CDFG of the following, larger example with three array reads. HereIn this embodiment, the latter method for producing the $START_{new}$ event is used.

```
x = a [i] ;  
10 y = a [j] ;  
z = a [k] ;
```

If several read sequences or read sequences and single read accesses occur for the same RAM, 1-events for detecting the *current accesses* must be generated for sequences of read
15 accesses. They are needed to separate the OUT-values relating to separate sequences. The ESEQ output just defined, sent through a 1-CONSTANT, ~~achieves~~may achieve this. It ~~is~~may
be again “N to 1 connected” to the other accesses’ START signals (for single accesses) or ESEQ outputs sent through 0-CONSTANT (for access sequences). The resulting event
ismay be used to control a first-stage DEMUX which is inserted to select the relevant OUT
20 output data packets of the sequence as described above for the basic method. ~~Refer to the second example (Figures 15 and 16) in Section 4.3 for a~~ complete example. is provided below under the heading “More Examples” with reference to Figs. 68 and 69.

4.2.8—Input and Output Ports

25 Input and output ports ~~are~~may be processed similar to vector accesses. A read from an input port is like an array read without an address. The input data packet ismay be sent to DEMUX functions which may send it to the correct subsequent operators. The STOP signal ismay be generated in the same way as described above for RAM accesses by combining the INPORT’s U output with the current and other START signals.

30 Output ports may control the data packets by GATEs like array write accesses. The STOP signal ismay also be created as for RAM accesses.

More Examples

Fig. 7.66 shows the generated CDFG for the following for loop.

```
    a = b+c;
    for (i=0; i<=10; i++) {
5      a = a + i ;
      x[i] = k;
    }
```

In this example, $IN1 = \{a\}$ and $IN2 = \{k\}$ (cf. (In this regard, see Fig. 10.18). The MERGE function for variable a is may be replaced by a 2:1 data connection as mentioned in the footnote of Section 4.2.5. Note above under the heading “FOR Loops.” It is noted that only one data packet arrives for variables b, c, and k, and one final packet is produced for a (out). ~~forbody~~ Forbody does not use a START event since both operations (the adder and the RAM write) are dataflow-controlled by the counter anyway. But the RAM’s EWR output is may be the forbody’s $START_{new}$ and may be connected to CNT’s NEXT input. Note It is noted that the pipelining optimization, cf. Section 4.2.6, (see that which is discussed under the heading “Vectorization and Pipelining”) was not applied here. If it is applied (which is possible for this loop), CNT’s NEXT input is not connected, cf. See Fig. 11.67. Here, the loop iterations overlap. $START_{new}$ is generated from CNT’s U output and forbody’s $START_{new}$, (i.e., RAM’s EWR output), as defined at the end of Section 4.2.5. the discussion under the heading “FOR Loops.”

The following program ~~contains~~ includes a vectorizable (pipelined) loop with one write access to array (RAM) x and a sequence of two read accesses to array (RAM) y. After the loop, another single read access to y occurs.

```
    z = 0;
    for (i=0; i<=10; i++) {
        x[i] = i;
        z = z + y[i] + y [2*i];
30    }
    a = y[k];
```

Fig. 15.68 shows the intermediate CDFG generated before the array access Phase 2 transformation is applied. The pipelined loop is may be controlled as follows: Within the

loop, separate START signals for write accesses to x and read accesses to y ~~are~~may be used. The reentry to the forbody ~~is~~may also be controlled by two ~~in-dependent~~independent signals (“cycle1” and “cycle2”). For the read accesses, “cycle2” ~~guarantees~~may guarantee that the read y accesses occur in the correct order. But the beginning of an iteration for read y and write x accesses is not synchronized. Only at loop exit all accesses must be finished, which ~~is~~may be guaranteed by signal “loop finished”. The single read access ~~is~~may be completely independent of the loop.

Fig. 16.69 shows the final CDFG after Phase 2. ~~Note~~It is noted that “cycle1” is removed since a single write access needs no additional control, and “cycle2” is removed since the inserted MERGE and DEMUX functions automatically guarantee the correct execution order. The read y accesses are not independent anymore since they all refer to the same RAM, and the functions have been merged. ESEQs have been allocated to control the MERGE and DEMUX functions of the read sequence, and for the first-stage DEMUX functions which separate the read OUT values for the read sequence and for the final single read access. The ECOMBs, 1-FILTERs, 0-CONSTANTS and 1-CONSTANTS are allocated as described in Section 4.2.7, Phase 2, with respect to Phase 2 under the heading “Array Accesses” to generate correct control events for the GATEs and DEMUX functions.

References

- [1] ~~A.V. Aho, R. Sethi, and J.D. Ullman. *Compilers—Principles, Techniques, and Tools*. Addison-Wesley, 1986.~~
- [2] ~~The Stanford SUIF Compiler Group. Homepage
5 <http://suif.stanford.edu>.~~
- [3] ~~A.H. Veen. Dataflow architecture. *ACM Computing Surveys*, 18(4), December 1986.~~
- [4] ~~S.J. Allan and A.E. Oldehoeft. A flow analysis procedure for the translation of high-level languages to a data flow language. *IEEE Transactions on Computers*, C-29(9):826–831, September 1980.~~
- 10 [5] ~~D. Galloway. The transmogrifier C hardware description language and compiler for FPGAs. In *Proc. FPGAs for Custom Computing Machines*, pages 136–144. IEEE Computer Society Press, 1995.~~
- [6] ~~M. Weinhardt and W. Luk. Pipeline vectorization. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 20(2), February 2001.~~
- 15

20

ABSTRACT

In a data-processing method, first result data may be obtained using a plurality of configurable coarse-granular elements, the first result data may be written into a memory that includes spatially separate first and second memory areas and that is connected via a bus to the plurality of configurable coarse-granular elements, the first result data may be subsequently read out from the memory, and the first result data may be subsequently processed using the plurality of configurable coarse-granular elements. In a first configuration, the first memory area may be configured as a write memory, and the second memory area may be configured as a read memory. Subsequent to writing to and reading from the memory in accordance with the first configuration, the first memory area may be configured as a read memory, and the second memory area may be configured as a write memory.